



Europäisches Patentamt  
European Patent Office  
Office européen des brevets



(11) EP 1 286 272 A2

(12)

## EUROPEAN PATENT APPLICATION

(43) Date of publication:  
26.02.2003 Bulletin 2003/09

(51) Int Cl.7: G06F 15/163

(21) Application number: 02255603.9

(22) Date of filing: 09.08.2002

(84) Designated Contracting States:  
AT BE BG CH CY CZ DE DK EE ES FI FR GB GR  
IE IT LI LU MC NL PT SE SK TR  
Designated Extension States:  
AL LT LV MK RO SI

(72) Inventors:  
• Garnett, Paul Jeffrey  
Camberley, Surrey GU15 3TS (GB)  
• King, James E  
Berkshire RG40 2JE (GB)  
• Mayhead, Martin P  
Surrey GU26 6QY (GB)  
• Heffernan, Peter  
Hampshire SO32 5AX (GB)

(30) Priority: 10.08.2001 US 311505 P  
10.08.2001 GB 0119558  
21.03.2002 US 366430 P  
21.03.2002 GB 0206719  
14.06.2002 US 172719  
14.06.2002 US 172721  
14.06.2002 US 172723

(74) Representative: Harris, Ian Richard et al  
D. Young & Co.,  
21 New Fetter Lane  
London EC4A 1DA (GB)

(71) Applicant: SUN MICROSYSTEMS, INC.  
Palo Alto, California 94303 (US)

(54) Interface standard support in modular computer systems

(57) An information processing module in the form of a blade server can communicate information signals with other modules of a modular computer system such as a high density server system under a selectable information protocol (for example Ethernet or Infiniband). The information processing module can include one or more information connections for communicating information signals. One or more serializer/deserializers

can be connected between the module processor(s) and the information connection(s). A selectable information protocol interface can be interconnected between the processors and the serializer/deserializer. A plurality of different information protocols can be used and be selectively connectable between the module processor(s) and the serializer/deserializer(s) using one or more multiplexer/demultiplexers.




Fig. 20A

## Description

### BACKGROUND OF THE INVENTION

[0001] The present invention relates to computer systems, in particular to multiprocessor systems, for example multiprocessor server systems.

[0002] One application for the present invention relates to high density computer systems, for example, computer server systems for telecommunications applications. In telecommunications applications, it is important to provide high reliability and high capacity of operation. Various approaches have been taken to providing such high-performance, high reliability systems. Typically such systems are designed around providing redundant resources so that if one component of the system develops a fault, the system remains operational using the redundant resources. Fault tolerance can also be achieved, for example, with multiprocessor systems that provide redundancy through dynamic, e.g., software-controlled, task distribution. High density systems are typically rack mountable, with one or more processor systems occupying a shelf in the rack. The trend in recent times is to make the computers with smaller form factors. This means that more computers can be located in a rack. This has the advantage of increasing the processing density within the racks, and also the advantage of reducing the distance between the computer systems.

[0003] The present invention relates to providing support for multiple interface standards in a modular computer system, for example, a high density computer system.

### SUMMARY OF THE INVENTION

[0004] An aspect of the present invention provides an information processing module for a modular computer system. The information processing module can include one or more processors operable to perform information processing. One or more (e.g., two) module information connections can be provided for communicating information signals. One or more serializer/deserializers can be connected between the processor(s) and the module information connection(s). One or more information protocol interfaces can be connected between the processor(s) and the serializer/deserializer(s).

[0005] The information processing module can be a server blade of a high density computer system that supports a plurality of server blades. By employing a serializer/deserializer for the interface between the processor(s) of the information processing module and other components of the modular computer system, a blade server system can be provided that can support a plurality of different interface protocols. An embodiment of the invention can provide an information processing module (e.g., a server blade) that can readily be inserted into a modular computing system operable under one

of a plurality of different information protocols.

[0006] To support a desired information protocol for a particular application, an appropriate type of information protocol interface (e.g., an Ethernet interface, or an Infiniband interface) can be selected and installed in the information processing module by supplying this to a customer. Alternatively, by providing a plurality of different types of information protocol interfaces for supporting different information protocols in the information processing module, the customer can select an information protocol to be used on initialization or in use of the information processing module.

[0007] Switching logic (e.g., a multiplexer/demultiplexer) can be employed to select an appropriate one of a plurality of information protocol interfaces provided by the information processing module. The information protocol interface can be, by way of example, an Ethernet interface or an Infiniband interface.

[0008] The module information signal interfaces can include one or more common module information protocol interfaces connected to one or more module physical layer interfaces. The, or each, common module information protocol interface can be connectable to one or more selectable module information connections via the module physical layer interface(s) (serializers/deserializers). The provision of a common information protocol interface that is selectively connectable to one or more of the module information connections enables the provision of redundant information connections at low cost.

[0009] One or more module management connections can be operable to enable communication of management signals with one or more service controllers being operable to perform system management functions via the module management connection(s). By providing an information signal interface between a module processor and a module information connection and providing a module service controller operating via a module management connection, flexibility on the choice of a processor and an internal structure for the information processing module is achieved. The information signal interface and the module service controller serve to provide a buffer between the information processing module and the remainder of the modular computer system. The information signal interface isolates the processor from the external information protocol used. The service controller enables the processor to be relieved of many of the system management functions, these being performed via the service controller and the dedicated management connection.

[0010] Another aspect of the invention provides a server blade including one or more processors, one or more (e.g., two) information connection for communicating information signals, one or more information protocol interfaces and one or more serializer/deserializers connected between the processor(s) and the information connection(s) for communicating information signals in accordance with a selectable information proto-

col.

[0011] A further aspect of the invention provides a modular computer system. The modular computer system (for example a high density server system) can include a carrier having one or more module receiving locations that are each configured to receive an information processing module (for example a server blade). The module receiving location can include one or more (e.g., two) carrier information connections operable to interconnect with a corresponding module information connection of a received information processing module. The carrier can further include one or more (e.g., two) switches connectable to a respective carrier information connection of the module receiving location(s). The switch or switches can connect the information processing module to an external information network in accordance with a plurality of information protocols.

[0012] The carrier can also include an interconnection member (e.g., a midplane) carrying information connectors and interconnection pathways, the interconnection pathways being configured to support a plurality of information protocols.

[0013] Another aspect of the invention provides a method of interconnecting an information processing module (for example a server blade) to other components of a modular computer system (for example a high density computer system) where the information processing module includes one or more processors operable to perform information processing. The method includes the processor(s) communicating information signals via one or more different information protocol interfaces and one or more serializer/deserializers to one or more (e.g., two) information connections.

[0014] The method can include performing a compatibility check as to whether an information processing module supports an information protocol employed for information communication between the information processing modules and other components of the modular computer system. The operation of the information processing module can be prevented or terminated (e.g., completion of initialization can be prevented) if it is determined that the information processing module does not support an information protocol employed for information communication between the information processing modules and other components of the modular computer system. It can involve selecting a compatible one of a plurality of information protocol interfaces provided by the information processing module.

[0015] A further aspect of the invention provides a method of interconnecting an information processing module to other components of a modular computer system wherein the information processing module includes one or more processor operable to perform information processing. The method includes the processor(s) communicating information signals over a plurality of information connections via one or more module information signal interfaces connected between the processor(s) and the module information connections.

The module information signal interface includes one or more common module information protocol interfaces connected to one or more module physical layer interfaces. The, or each, common module information protocol interface is connectable to one or more selectable module information connections via the module physical layer interface(s).

[0016] Another aspect of the invention provides a method of providing a processor technology independent information server blade for a modular server system. The method includes one or more blade processors communicating information signals with other modules of the server system via at least one blade information signal interface between the processor(s) and one or more blade information connections for communicating information signals. It further includes one or more service controllers in the blade performing system management functions within the blade and communicating management signals to other modules of the server system via one or more module management connections.

[0017] Further aspects and advantages of the invention will become apparent from the following description of particular embodiments.

## BRIEF DESCRIPTION OF THE DRAWINGS

[0018] Embodiments of the present invention will be described hereinafter, by way of example only, with reference to the accompanying drawings in which like reference signs relate to like elements and in which:

Figure 1 is a schematic representation of an architecture of a multiprocessor system for supporting a web site;

Figure 2 is a schematic representation of a racking system incorporating an example of a carrier in the form of a rack-mountable shelf according to a first example;

Figure 3 is a front view of an example of a carrier in the form of a shelf of Figure 2;

Figure 4 is a rear view of an example of the shelf of Figure 2;

Figures 5A, 5B and 5C are schematic perspective views and a schematic exploded view respectively of an example of an information processing cartridge for mounting in the shelf of Figure 2;

Figure 5D is a schematic perspective view of an example of an information processing cartridge for mounting in the shelf of Figure 2;

Figure 6 is a schematic perspective view of an example of a combined switch and service processor module for mounting in the shelf of Figure 2;

Figure 7 is a schematic perspective view of an example of a power supply module for mounting in the shelf of Figure 2;

Figures 8A, 8B and 8C are a schematic plan view and schematic perspective views, respectively, of an example of the chassis and midplane of the shelf of Figure 2;

Figure 8D is a schematic perspective view of a part of a base member of the shelf of Figure 2;

Figures 9A, 9B and 9C are schematic front, top and rear views, respectively, of an example of a midplane of the shelf of Figure 2;

Figure 10 is a schematic cross section view through the shelf of Figure 2;

Figure 11 is a functional block diagram of an example of an information processing subsystem for the information processing cartridge of Figure 5;

Figure 12 is a functional block diagram of an example of an information processing subsystem for the combined switch and service processor module of Figure 6;

Figure 13 is a functional block diagram of an example of a subsystem for the power supply unit of Figure 7;

Figure 14 is a functional block diagram showing the connectivity between the components of the shelf of Figure 2;

Figure 15 is a functional block diagram showing the external connectivity of the shelf of Figure 2;

Figure 16 is a schematic representation of a shelf showing the external connections from the shelf of Figure 2;

Figure 17 is a schematic representation of a rack mounted system comprising a plurality of such shelves;

Figure 18 is a functional block diagram of a further example of an information processing subsystem for the information processing cartridge of Figure 5;

Figures 19A and 19B illustrate different examples of information protocol interfaces for an information processing cartridge;

Figures 20A and 20B illustrate further examples of information protocol interfaces for an information processing cartridge;

Figures 21A and 21B illustrate further examples of an information protocol interface for an information processing cartridge; and

Figures 22A and 22B represents further examples of an information protocol interface for an information processing cartridge.

[0019] While the invention is susceptible to various modifications and alternative forms, specific embodiments are shown by way of example in the drawings and are herein described in detail. It should be understood, however, that drawings and detailed description thereto are not intended to limit the invention to the particular form disclosed, but on the contrary, the invention is to cover all modifications, equivalents and alternatives falling within the spirit and scope of the present invention as defined by the appended claims.

## DESCRIPTION OF PARTICULAR EMBODIMENTS

[0020] Embodiments and examples are described hereafter by way of example only in the following with reference to the accompanying drawings.

[0021] Shown in Figure 1 is an example of an application of a high capacity multiserver system 1 for implementing a network-connected web site such as, for example, an airline reservation system on the World Wide Web.

[0022] As shown in Figure 1, an external network 3 (e.g., the Internet) for communicating with a user 5 can be connected to gateways 7 which can be connected to an entry edge server group 9 implemented by a web farm as shown in Figure 1. The entry edge server group 9 forms an interface to the external network 3. The entry edge server group 9 can then be connected by switches 11 and a firewall 13 to a web edge server group 15 that can also be implemented as a web farm as shown in Figure 1. The web edge server group 15 can serve to cache web pages that are readily accessible to users 5 accessing the system 1 from the external network 3, for example for checking flight times, etc. The web edge server group can comprise a number of blade server (BS) shelves and a number of network addressable storage (NAS) shelves for storing critical data. The web edge server group 15 can be further connected by a further firewall 17 to a plurality of application servers 19, which can be responsible for, for example, processing flight reservations. The application servers 19 can then be connected via a further firewall 21 to computer systems 23, 25, for example, e-commerce services including financial services for receiving and processing payment for airline reservations.

[0023] As will be appreciated, the server system described above with reference to Figure 1 is only an example of a possible application for a multiprocessor server system. Multiprocessor server systems have many different applications and the present system is

not limited to being applicable for use in only one or a limited number of such applications, rather multiprocessor server systems as described herein are operable for use in many different applications. A non-exhaustive list of such alternative applications includes: e-commerce web server systems; telecommunications network server systems; LAN application and file server systems and remote vehicle control systems.

[0024] With reference to Figure 2, there is shown a schematic perspective representation of a rack system 31 as viewed from the front including left and right front uprights 32 and 33 and left and right rear uprights 34 and 35. The uprights can be formed with apertures for receiving shelf fixings (e.g., screws, bolts, clips, etc., for mounting brackets, slides, rails, etc.).

[0025] Also shown in Figure 2 is an example of a blade server shelf 41 mounted in the rack system 31. The shelf 41 forms a carrier configured to carry a plurality of information processing cartridges 43 located side by side along the shelf.

[0026] The term "shelf" is used herein in a conventional way to describe a structure that is mountable in rack system 31 and is configured to carry one or more components to form at least a part of a rack-mountable system. In the present example, the shelf 41 is three-dimensional, having a height (H), width (W) and depth (D). In the present example, one dimension (hereinafter described as the height, H) is smaller than the other dimensions (hereinafter described as the depth, D, and the width, W) to facilitate mounting of the shelf within the rack system 31. It will be appreciated that although the width and depth are typically constrained by the dimensions of the racking system for which the shelf is designed, there is more freedom as regard the height, subject to taking account of any appropriate standards and packaging considerations.

[0027] Each of the information processing cartridges contains at least one processor. Each information processing cartridge in the present example is operable as a server. In the described examples, the information processing cartridges are configured as robust enclosed modules.

[0028] In the example to be described in more detail presently, the information processing cartridges, when aligned in the carrier shelf, look like rectangular slabs, or blades. Accordingly, an information processing cartridge can be described as a blade. The information processing cartridges 43 comprise information processing modules enclosed in an enclosure, or housing, so that the information processing modules have the form of cartridges. Also, as the information processing cartridges are to operate as computer servers in the example described in more detail presently, an information processing cartridge 43 can also be described as a server blade. Accordingly, in the context of this example, the terms module, cartridge and blade are used interchangeably.

[0029] The illustrated example of a shelf 41 is config-

ured to carry sixteen information processing cartridges 43, each of which is removably mountable in a respective opening 45 in the front of the shelf, whereby the information processing cartridges can be inserted into and removed from the front of the shelf 41 without removing the shelf 41 from the rack system 31.

[0030] In the present example, the shelf 41 comprises a three-dimensional, generally rectangular, enclosure, or housing, 47 that is suitable for mounting in generic racking systems including both 4-post and 2-post systems. It can be mounted on fixed rigid rack mounting ears and/or a simple slide/support system. The present example is designed for standard 19"-wide racking (1"=25.4mm) as defined, for example, in the well-known IEC297 and EIA 310 specification standards with height corresponding to the so-called 3U (3 standard unit) height. For mounting such a 3U unit in such a 19"-wide racking system, with a depth of, say 25" or 30", the enclosure can be arranged with a height of up to about 130.5mm, a width of up to about 445 mm and a depth, including all hardware and fascias, but excluding cable management, of up to about 635mm, with the depth from the front-most point of a fascia to a rear I/O connector panel of a rear mounted Field Replaceable Unit (FRU) of about 610mm. Of course, other examples designed for other racking systems could have different dimensions.

[0031] This example of a shelf 41 has a single enclosure, or housing, 47 that houses a number of modular units or subsystems, the majority of which are replaceable in the field and are therefore known as Field Replaceable Units (FRUs). These modular units include the information processing cartridges 43.

[0032] The shelf enclosure 47 can be fabricated from sheet material (e.g., from steel sheet) to form a chassis portion 49 that includes a base 51, two sides 53 and 55, a front 57 and a rear 59. The word "front" as used here is merely used as a label herein to refer to the face, or wall 57 of the enclosure that is located at the main access side of the rack system 31 in use when the shelf is mounted therein. Similarly, the words "rear" and "side" are merely used as labels herein to refer to the faces, or walls 59, 53 and 55 that, in use, are located at those respective positions when the shelf is mounted in the rack system 31.

[0033] The openings 45 can be formed in the front face 57 for receiving the information processing cartridges 43 and, as will be explained later, apertures can also be formed in the rear face 59 for receiving further FRUs. The enclosure can further include a removable top cover 61 that can be secured to the chassis portion 49 by suitable fastening (e.g., screws). The apertures in the front and rear faces 57 and 59 allow at least some of the FRUs to be inserted into and/or removed from the shelf enclosure 47 via the front or the rear thereof, as appropriate, without removing the shelf from the racking. Access to components mounted in the shelf that are not accessible via one of the apertures in the front 47

and rear 59 faces can be achieved by removing the shelf enclosure 47 from the racking system 31 and then removing the top cover 61 of the shelf enclosure 47.

[0034] Figure 3 is a front view of an example of a shelf 41 for a first example. A plastic front bezel 63 can be provided that fits on the front face 57 (shown in Figure 2) of the chassis 49 of the shelf enclosure 47. The front bezel 63 can be formed as a unitary removable part that spans the whole width and height of the front of the shelf enclosure 47. The front bezel 63 could alternatively include a number of separate components, or mouldings. The front bezel can include a peripheral portion 64 that can provide areas for corporate and product branding marks, for identification and numbering for the information processing cartridge and for a bar code label (all not shown). One or more apertures 65 can be formed in the peripheral portion 64 of the bezel 63. The apertures 65 in the bezel can be arranged to align with one or more apertures (e.g. a slot (not shown in Figure 3) in the front face of the chassis. In use, air can pass through the apertures 65 to flow into the shelf enclosure 47 to reach FRUs that are mounted in the shelf enclosure 47 through the rear face 59 thereof. Air flowing through the aperture 65 flows into a plenum chamber 66 (not shown in Figure 3) to flow past the processing cartridges 43 to reach rear mounted FRUs. A central area 67 of the front bezel 63 can be open allowing access to the openings 45 in the front face 57 of the shelf enclosure 47 for insertion and removal of the information processing cartridges 43. Where no active module is mounted in a location for an information processing module, a blanking panel, or filler panel, such as the blanking panel 44, can be located in that location. LED indicators 69 can be mounted on a system indicator printed circuit board (not shown) behind a designated area of the bezel to provide an indication of system status via light guides incorporated into the bezel. A further system indicator board (also not shown) carrying LED indicators can be provided inside the shelf enclosure to be visible from the rear thereof.

[0035] As mentioned above, in the present example of a shelf, up to sixteen information processing cartridges 43 can be installed in respective openings 45 in the front face 57 thereof. The number of information processing cartridges 43 actually installed in any installation is dependent upon the system configuration required. Various features relating to the information processing cartridges 43 that are shown in Figure 3 will be described later.

[0036] Figure 4 illustrates the rear of the shelf unit of Figures 2 and 3. This shows two different types of FRU 71 and 81 (4 units in total) that have been inserted into respective apertures 72 and 82 in the rear of the shelf enclosure 47. The FRUs shown in Figure 4 include two Combined Switch and Service Processors (CSSPs) 71 and two Power Supply Units (PSUs) 81. Various features shown in Figure 4 will be described later.

[0037] Before proceeding with a more detailed de-

scription of each of the FRUs 43, 71 and 81 introduced so far and of the construction of the shelf 41, there follows a brief description of an information processing cartridge 43, a CSSP 71 and a PSU 81 with reference to Figures 3, 4, 5, 6 and 7.

[0038] Figure 5A provides a perspective view, partly from the rear, of an information processing cartridge 43. Figure 5B provides a perspective view, partly from the front of the same information processing cartridge 43. Figure 5C provides an exploded perspective view of the construction of the information processing cartridge 43. Here it is to be noted that the term "rear" is applied in the context of the position, when installed, of the information processing cartridge, with respect to the shelf 41 (i.e. in this case the "rear" of the information processing cartridge 43 is the innermost part of the information processing cartridge when it is inserted in the shelf 41). Likewise, "front" refers in the present context to the outermost part of the information processing cartridge when it is inserted in the shelf 41.

[0039] With reference to Figures 3, 5A and 5B, it will be noted that information processing cartridges are three-dimensional, having a height (h), width (w) and depth (d). If, as in the present example, the information processing cartridges are to be arranged in a one-dimensional array (a row) across the shelf, then efficient packing for the information processing cartridges is achieved where one dimension (here the width, w) is smaller than the other dimensions (here the depth, d, and the height, h). In a particular example, the enclosure of a processing cartridge 43 has a height h, width w and depth d of 115mm, 26mm and 315mm, respectively, although of course, other examples may have different dimensions.

[0040] It will be noted that an enclosure 101 of the present example of an information processing cartridge 43 has six, generally rectangular, faces. For convenience only, the face that is visible from the front of the racking when an information processing cartridge 43 is mounted in the shelf 41 is known as the front face 102. The opposite face is known as the rear face 103. In the present example these two faces, as well as top and bottom faces 104 and 105, and the side faces 106 and 107 have the shape of elongate rectangles.

[0041] Although in the present example the information processing cartridges have six generally rectangular faces, it will be appreciated that other examples could have other configurations. For example, rather than having a generally rectangular shape, the side faces of another example of an information processing cartridge could have the general shape of a triangle (whereby the information processing cartridge may then only have five faces), a pentagon (whereby the information processing cartridge may then have seven faces), and so on. Indeed, one or more or all of the edges could be curved. However, it will be appreciated that the present configuration provides advantages for example, in terms of manufacturing, engineering and packing den-

sity within a shelf 41.

[0042] In this example, the information processing cartridge enclosure 101 is fabricated from two housing portions. The first housing portion 109 can be fabricated from, for example, a metal (e.g., pressed steel) and can include one side face 107, the rear face 103 and part of the top face 104. The part of the top face formed by the metal portion 109 is given the reference 1041 in Figure 5A. The second housing portion 110 can be fabricated from, for example, a plastics material and can include the other side face 108 and the front faces 102 and bottom face 105 and the remaining part 1042 of the top face 104. In the present example, a plastics material used is PolyCarbonate Acrylonitrile Butadiene Styrene (PCABS), however many other plastics materials such as other ABSs and nylons may be used. The metal portion 109 also includes a part 1021 corresponding to the front face 102 and a part 1051 corresponding to a part of the bottom face 105. The parts are covered by the plastics portion 110 when enclosure 101 is fully assembled. These parts are visible in the exploded view of Figure 5C. The two housing portions 109 and 110 are secured to one another by fixings such as screws 118.

[0043] As shown in Figure 5A, a groove 108 runs along the length of the plastic top face portion 1042. This groove 108 is provided for interfacing with a guide member of the shelf chassis (not shown in Figure 5A, but see Figures 8B and 8C). A similar groove 113 is formed in the bottom face 105 of the plastics portion 110 as shown in Figures 5B and 5C. This groove 113 is provided for interfacing with a guide member of the shelf chassis (not shown in Figures 5B and 5C, but see Figure 8D).

[0044] The provision of the enclosure 101 means that the information processing cartridge 43 can safely be handled by an operator who is not a skilled technician. Also, through the use of the enclosure 101, the information processing cartridge is a robust unit that protects its inner workings from the outside environment and vice versa. The use of a conductive enclosure, e.g., a metal enclosure, can result in the information processing cartridge including its own electromagnetic shielding. To this end, the enclosure 101 can also be provided with an internal electromagnetic interference (EMI) shielding member 116 as shown in Figure 5C. The EMI shielding member can be located between the internal components 112 of the information processing cartridge 43 (not described in detail with reference to Figures 5A to 5C, but see Figures 11 and 18 below) and the plastics portion 110. The EMI shielding member can be secured to the plastics portion 110, for example by heat bonding or an adhesive. In other examples, the plastics portion could have a conductive layer deposited thereon, or conductive plastics material could be used. In this example, the EMI shielding member 116 is provided with electromagnetic interference (EMI) fingers 114 to ensure good contact with the shelf chassis and the adjacent components. These fingers 114 extend through EMI finger apertures 1141 in the plastics portion 110.

[0045] As shown in Figures 3, 5B and 5C, the information processing cartridge 43 incorporates an injector/ejector handle 111 on the front face 102 for facilitating insertion and latching of the information processing cartridge 43 within an aperture in the shelf 41.

[0046] As shown in those Figures, the handle 111 of the injector/ejector lever extends substantially the whole height of the front face of the information processing cartridge 43, thereby increasing the mechanical advantage and facilitating injection and ejection of the information processing cartridge 43. As is further shown in those Figures, the front face 102 of the information processing cartridge 43 has perforations 115, in the present example slits, to allow for airflow into the information processing cartridge 43. The front face part 1021 of the metal portion 109 has perforations corresponding to those in the front face 102 such that airflow into the information processing cartridge 43 is not impeded. It will be noted in Figures 3, 5B and 5C that the handle 111 is narrower in its middle than at its ends. This reduces any masking effect of the handle 111 on the airflow to the perforations 115 in the front face of the information processing cartridge 43 and facilitates grasping of handle 111. The handle 111 can be bowed out from the front face of the information processing cartridge to further facilitate grasping thereof and to reduce any masking effect with regard to the perforations. In other examples, the handle 111 could have an open frame-like structure to further facilitate airflow.

[0047] As shown in Figure 5A, the rear face 103 of the information processing cartridge 43 also has perforations 117 to allow for air to be exhausted from the rear of the information processing cartridge 43. A fan can be located within the enclosure 101 of an information processing cartridge 43. In the present example the fan is a combined radial fan (blower) and heat sink to direct cooling air onto a processor of the information processing cartridge, which fan also encourages air movement within the enclosure 101. LED indicators 119 (see Figure 3 and 5B) can be provided on the front face 102 of an information processing cartridge to indicate whether power is on, whether service intervention is required and whether the information processing cartridge 43 can be removed. A connector 120, for example a 40 way single connector attachment (SCA-2) connector (a small computer systems interface (SCSI) connector), can be provided at the rear of the information processing cartridge 43 for electrical connection of the information processing cartridge 43 within the shelf 41. The connector 120 is advantageously able to withstand repeated removals and insertions of the enclosure 101 from and into a shelf 41. The connector arrangement can include a guide pin arrangement to prevent module misalignment during insertion of the information processing cartridge into the receiving location.

[0048] Thus an example of features and the construction of an information processing cartridge enclosure has been described. Although particular materials and

constructions have been described, it will be appreciated that other examples could be employed. Indeed, it will be appreciated that this particular example relates to only a possible form for the processing cartridge 43. An example of an alternative construction will now be described with reference to Figure 5D.

[0049] Figure 5D provides a perspective view, partly from the rear, of an alternative information processing cartridge 43. Here it is to be noted that the term "rear" is applied in the context of the position, when installed, of the information processing cartridge, with respect to the shelf 41 (i.e. in this case the "rear" of the information processing cartridge 43 is the innermost part of the information processing cartridge when it is inserted in the shelf 41).

[0050] In this example, the information processing cartridge enclosure 101 is fabricated from pressed steel to form two chassis portions. The first portion 234 includes one side face 107, and part of each of the front and rear faces 102 and 103 and the top and bottom faces 104 and 105. The second portion 235 includes the other side face 108 and the remaining part of each of the front and rear faces 102 and 103 and the top and bottom faces 104 and 105. The two chassis portions 234 and 235 meet at a groove 236 and are secured to one another by fixings (e.g., one or more screws, not shown). Grooves 236 run along the top and bottom faces 104 and 105 of the enclosure 101 and are provided for interfacing with guide rails of the shelf chassis 49 (not shown in Figure 5, but see Figure 8C). A cover portion that is secured to the chassis portion forms the other side face 106. It will be appreciated however, that in another example, the chassis portions could be joined at a position other than the groove 108, with the groove 108 being formed entirely in one of the chassis portions. Alternatively, the enclosure 101 may be constructed from a number of sheets of steel, with each sheet forming one of the faces.

[0051] Figure 6 provides a perspective view, partly from the front, of a Combined Switch and Service Processor (CSSP) cartridge (or CSSP) 71. Here it is to be noted that the term "front" is applied in the context of the position, when installed, of the CSSP cartridge 71, with respect to the shelf 41 (i.e. in this case the "front" of the CSSP cartridge 71 is the innermost part of the CSSP cartridge 71 when it is inserted in the shelf 41).

[0052] With reference to Figures 4 and 6, it will be noted that a CSSP cartridge 71 is three-dimensional, having a height (h), width (w) and depth (d). In a particular example, the enclosure of a CSSP 71 has a height h, width w and depth d of 43mm, 202mm and 278mm, respectively, although of course, other examples may have different dimensions.

[0053] An enclosure 121 of present example of a CSSP cartridge 71 has six, generally rectangular, faces. For convenience only, the face that is visible from the rear of the racking when a CSSP cartridge 71 is mounted in the shelf 41 is known as the rear face 122. The

opposite face is known as the front face 123. In the present example these two faces, as well as side faces 126 and 127 have the shape of elongate rectangles. The top and bottom faces 124 and 125 are also rectangular, but not elongate in the manner of the front, rear, top and bottom faces. Although in the present example the CSSP cartridges have six generally rectangular faces, as for the information processing cartridges 43 it will be appreciated that other examples could have other configurations.

[0054] In this example, the CSSP enclosure 121 is fabricated from steel sheet to form a chassis portion that includes the bottom face 125, the front and rear faces 122 and 123 and the side faces 126 and 127. A cover portion that is secured to the chassis portion forms the other top face 124. The cover portion is secured to the chassis portion by suitable fixings, for example one or more screws 128. It will be appreciated however, that in another example, other faces, or portions, of the enclosure could form the chassis and the cover portions. The provision of the enclosure 121 means that the CSSP cartridge 71 can safely be handled by an operator who is not a skilled technician. Also, through the use of the enclosure 121, the switch cartridge is a robust unit that protects its inner workings from the outside environment and vice versa. The use of a conductive enclosure, e.g., a metal enclosure, means that the CSSP cartridge includes its own electromagnetic shielding. To this end the CSSP enclosure 121 is provided with EMI fingers 129 to ensure good contact with the shelf chassis and the adjacent components.

[0055] As shown in Figure 4, the CSSP cartridge 71 incorporates two D-shaped handles 132 to facilitate insertion and removal of the CSSP cartridge 71 with respect to an aperture 72 in the rear face of the shelf enclosure. A latch member 131 can be pivotably mounted on a plate that can be secured (e.g., using screws) to the rear face of the shelf enclosure. The latch member 131 is configured to engage one of the handles 132 and to secure the CSSP cartridge 71 in place. In other examples, the CSSP cartridge 71 could be provided with an injector/ejector handle in a manner similar to the information processing cartridge. As shown in Figure 6, the front face 123 of the CSSP cartridge 71 has perforations 133 to allow for airflow into the CSSP cartridge 71. As shown in Figure 4, the rear face 122 of the CSSP cartridge 71 has perforations 135 to allow for air to be exhausted from the rear of the CSSP cartridge 71.

[0056] At least one fan can be located, for example behind the perforated portion 135 of the rear face, in a CSSP cartridge 71 to channel cooling air through the CSSP cartridge 71 from the front to the rear. In this particular example shown, two fans are provided, one behind each set of perforations 135. LED indicators 137, as shown in Figure 4, can be provided on the rear face 122 of the CSSP enclosure 121 to indicate whether power is on, whether service intervention is required and whether the switch can be removed. Additional link sta-



tus indicators can be provided integral to 2x4 stacked RJ-45 connectors 139, also shown in Figure 4. As shown in Figure 4, electrical connections 141 can be provided at the front face of the CSSP (i.e. on the face that in use is inside the shelf enclosure 47). Suitable connections for use in the present example include a connector for power connections, a connector for serial management data connections and a connector for information connections. In the present example, information connections are implemented using an Ethernet information communication protocol, e.g. at 1 Gigabit (Gb). However other protocols could equally be used, for example the Infiniband information communication protocol. The connector arrangement can include a guide pin arrangement to prevent module misalignment during insertion of the CSSP module into the receiving location. For this purpose, guide pin holes 142 can be provided on the front face 123 into which guide pins may pass to aid module alignment.

[0057] In the present example, up to two CSSPs 71 can be mounted at any one time at the rear of the shelf unit in corresponding apertures 72 in the rear face of the shelf enclosure 47. The number of CSSPs 71 provided in any particular implementation depends upon system configuration and the need, or otherwise, for redundancy.

[0058] It will be appreciated that one possible construction of the CSSP cartridge 71 has been described and that as for the information processing cartridge 43, other examples could employ other materials and/or constructions.

[0059] Figure 7 provides a perspective view, partly from the front, of a power supply unit (PSU) cartridge 81. Here it is to be noted that the term "front" is applied in the context of the position, when installed, of the PSU cartridge 81, with respect to the shelf 41 (i.e. in this case the "front" of the PSU cartridge 81 is the innermost part of the PSU cartridge 81 when it is inserted in the shelf 41).

[0060] With reference to Figures 4 and 7, it will be noted that PSU cartridge 81 is three-dimensional, having a height (h), width (w) and depth (d). In this particular example, the order to provide for a dense packing of the FRUs in the rear of the shelf 41, the PSU cartridge 81 has two dimensions (hereinafter described as the width, w, and the depth, d) that are generally similar. In a particular example, the enclosure of a PSU cartridge 81 has a height h, width w and depth d of 83mm, 202mm and 276mm, respectively, although of course, other examples may have different dimensions.

[0061] An enclosure 145 of present example of a PSU cartridge 81 is of generally oblong shape, but has the "top" "front" edge cut away to form an additional "top" "front" sloping face. The enclosure 145 therefore has five, generally rectangular, faces and two faces of generally rectangular shape with one corner cut away. For convenience only, the face that is visible from the rear of the racking when the PSU cartridge 81 is mounted in

the shelf 41 is known as the rear face 146. The opposite face is known as the front face 147. In the present example these two faces and the two side faces 150, 151 are of elongate, generally rectangular shape with one corner cut away, given that the width and depth of the PSU cartridge are similar, whereas the top and bottom faces 148, 149, although still rectangular, are not, in this example, notably elongate. A top front face 148a is present at the top front of the enclosure. Thus the front of the enclosure is sloped at the top edge. As for the information processing cartridges 43, however, it will be appreciated that other examples could have other configurations.

[0062] In this example, the PSU cartridge enclosure 145 is fabricated from steel sheet to form a housing portion that includes the bottom face 149, the side faces 150 and 151 and the front and rear faces 146 and 147. Cover portions that are secured to the housing portion form the top face 148 and top front face 148a. The cover portions are secured to the chassis portion by suitable fixings, for example one or more screws 152. It will be appreciated however, that in another example, other faces, or portions, of the enclosure could form the chassis and the cover portions. The provision of the enclosure 145 means that the PSU cartridge 81 can safely be handled by an operator who is not a skilled technician. Also, through the use of the enclosure 145, the PSU cartridge 81 is a robust unit that protects its inner workings from the outside environment and vice versa. The use of a conductive enclosure, e.g., a metal enclosure, means that the PSU cartridge includes its own electromagnetic shielding. To this end the PSU enclosure 145 is provided with EMI fingers 153 to ensure good contact with the shelf chassis and the adjacent components.

[0063] As shown in Figure 4, the PSU cartridge 81 incorporates two D-shaped handles 156 to facilitate insertion and removal of the PSU cartridge 81 with respect to an aperture 82 in the rear face of the shelf enclosure. A latch member 155 can be pivotably mounted on a plate that can be secured (e.g., using screws) to the rear face of the shelf enclosure. The latch member 155 is configured to engage one of the handles 156 and to secure the PSU 81 in place. In other examples, the PSU 81 could be provided with an injector/ejector handle in a manner similar to the information processing cartridge. As shown in Figure 7, the front face 147 of the PSU cartridge 81 has perforations 157 to allow for air-flow into the PSU cartridge 81. As shown in Figure 4, the rear face 146 of the PSU cartridge 81 also has perforations 159 to allow for air to be exhausted from the rear of the PSU cartridge 81.

[0064] A pair of fans can be located behind the perforated portions 159 of the rear face of a PSU cartridge 81 to channel cooling air through the PSU cartridge from the front to the rear. LED indicators 161 can be provided on the rear face 146 of the PSU enclosure 81 to indicate whether input power is good, whether output power is

good, whether service intervention is required and whether the PSU can be removed. Electrical connectors 163 can be provided at the front face of the PSU (i.e. on the face that in use is inside the shelf enclosure 47) for connection to the shelf. The PSU 81 of the present example may suitably employ an SSI-MPS (Server Systems Interface-Midrange Power Supply) compliant right angle connector at the front face 147 of the PSU 81 to connect to the shelf 41. The power inlet 83 for each PSU 81 can incorporate a cable/connector retention mechanism (not shown) on the rear face 146 of the PSU to prevent accidental or malicious removal of the power input cord from the PSU 81.

[0065] In the present example, the shelf unit enclosure provides slots at the rear of the shelf unit for two hot-swappable, AC input PSUs 81. To provide redundant power supply, both PSU slots are populated. More generally, N+M power supplies can be provided, where N is the minimum number of power supply units needed to support the components mounted in the shelf, and M is a selectable number of 0, 1 or more power supplies to provide for redundancy. In the present example, N = M = 1.

[0066] It will be appreciated that one possible construction of the CSSP cartridge 71 has been described and that as for the information processing cartridge 43, other examples could employ other materials and/or constructions.

[0067] If the full complement of information processing cartridges or switches are not fitted to the enclosure, then blanking panels/modules (e.g., the blanking panels 44 shown in Figure 3) are fitted to all unpopulated positions to maintain the correct airflow and thermal characteristics of the shelf, a safe internal operating temperature for the shelf life expectancy of all shelf components, electromagnetic compliance (EMC) containment and electrostatic discharge (ESD) containment.

[0068] It will be noted that each of the FRUs, such as the information processing cartridges 43, is advantageously contained in its own robust enclosure to facilitate EMC containment, ESD containment, handling, storage and transportation. Each FRU can be configured as a 'sealed' unit in the sense that it can be configured not to have field or customer serviceable parts internally. The FRUs can be configured readily to plug into the shelf enclosure and to be hot swappable. The FRUs can be keyed to prevent incorrect positioning and insertion into the shelf enclosure and are arranged positively to be retained in the shelf by a latching/locking mechanism.

[0069] The examples of FRUs described above are not provided with removable media. In the present example, internal data storage is provided by 2.5" IDE 9.5mm or 12.7mm profile hard disk drive (HDD) devices mounted internally in each information processing cartridge 43 and in the CSSP cartridge 71. The drives are not considered as FRUs and are not hot-swappable disk drives in the present example, although they could be

in other examples. In other examples, the information processing cartridges can be configured without internal hard disk drives.

[0070] The internal configuration of the shelf 41 and a midplane 171 contained therein is described in the following with reference to Figures 8A, 8B 8C and 8D, and Figures 9A, 9B and 9C.

[0071] Figure 8A is a schematic plan view showing the internal configuration of an example of a shelf 41 with the cover 61 removed. Figure 8B is a schematic perspective view from above the rear of the chassis portion 47 of the shelf enclosure with the field replaceable units removed. Figure 8C is a schematic perspective view from below the front of the chassis portion 47 of the shelf enclosure with the field replaceable units and the base 51 removed. Figure 8D is a schematic perspective view from the front and above a part of the base 51 of the shelf 41. Figures 9A, 9B and 9C are, respectively, front, top and rear views of the midplane 171. In this example, the midplane is, in use, mounted vertically within the shelf 41 extending across the width W of the shelf 41 at a position approximately half way between the front and the rear of the shelf 41.

[0072] The vertically mounted midplane 171 extends, in this example, across the shelf 41 and allows for the electrical interconnection of the FRUs. The various apertures in the front and rear faces 57 and 59 of the shelf 41, in combination with the midplane 171, can be provided with guides (e.g., rails 181) and keying e.g., offset connector positioning for the insertion of the FRUs into the enclosure and midplane 171. The midplane 171 can be a double-sided, or multi-layer printed circuit board (PCB) assembly that can be mounted vertically in a rigid manner within the enclosure. It can carry connectors 175 on a front surface 172 for making electrical connection with corresponding connectors 120 on the information processing cartridges 43. It can also carry connectors 177 and 179 on rear surface 173 for making electrical connection with corresponding connectors 141 and 163 on the CSSPs 71 and the PSUs 81, respectively. Conductive tracks (not shown) on and through the midplane 171 can be provided to interconnect the various connectors. In addition, the midplane can provide connectors for receiving corresponding connectors connected to first and second indicator boards 183 and 184 that each carry a respective set of LED indicators 69. In the present example, the midplane 171 is not configured as a FRU and is not hot swappable. It is perforated to facilitate airflow through the shelf 41. The midplane 171 can include openings 185, which cooperate with openings in the enclosures of the FRUs 43 and 81, to provide a path for cooling air to pass from the front to the rear of the shelf 41, the cooling air being driven by fans in one or more of the FRUs, for example in the PSUs 81, possibly also in the information processing cartridges 43.

[0073] A plenum chamber floor member 94 can extend horizontally from the front of the midplane 171 to

the front face 57 of the shelf enclosure, or chassis 47. The member 94 provides a floor for a plenum chamber 66, which is supplied with air via the apertures 65 in the front bezel and, in the illustrated example, the slot shaped aperture 68 in the front face 57 of the shelf enclosure 47. Although, for reasons of ease of illustration a slot shaped aperture 68 is shown, a plurality of apertures 68 aligned with the blade receiving locations may be provided. The aperture or apertures 68 can serve both as air vents for a flow of air to the plenum chamber 66, and also as latching locations for latching portions at the top of the injector/ejector levers 111 for the blades shown in Figures 5B and 5C. The top and sides of the plenum chamber are provided by the top cover 61 and side faces 53 and 54 of the shelf enclosure 47.

[0074] A plurality of cartridge guides 97 can be provided at the underside of the plenum chamber floor member 94. In the present example, these guides comprise sprung wire members, e.g., of a resilient metal such as spring steel, that are attached to the top surface of the plenum chamber floor member 94 and extend through a plurality of apertures therethrough to result in a row of guides 97 at the underside of the plenum chamber floor member 94. This arrangement is shown in Figures 8B and 8C. In Figure 8B, the sprung wire members 98 are shown attached to the top surface of the plenum chamber floor member 94. In the present example, the sprung wire members 98 are arranged in pairs, such that two guides 97 are provided by each spring clip 98. In Figure 8C, the guides 97 formed by the protruding portions of the sprung wire members 98 are shown at the underside of the plenum chamber floor member 94. Each guide 97 is advantageously positioned so as to interface with the groove 108 in the plastics material in the upper face 104 of a processing cartridge 43 as shown in Figures 5A-5C to aid correct alignment and to facilitate insertion of the processing cartridge during insertion of the cartridge into the shelf 41. The use of the spring clip as a guide 97 also serves to urge the processing cartridge downwards to provide a secure mounting of the processing cartridge 43, to take account of manufacturing and operational tolerances and to assist in insertion of the processing cartridge where an operator does not align this absolutely correctly.

[0075] A further row of cartridge guides 99 can be provided at the upper surface of the base 51 of the shelf 41. In the present example, as shown in Figure 8D, these guides 99 have a rail like form, which can be achieved by punching or stamping through the base 51 of the shelf 41. In this example each guide, or rail, 99 includes a pair of upstands separated by an aperture 100 through the base 51. The size of the aperture 100 can correspond to the width between the upstands. The separation of the upstands is selected so that the overall width of the resulting rails is slightly less than the width of a groove formed in the lower face of an information processing cartridge 43. Thus, each guide 97 is advantageously arranged so as to interface with the groove

1113 in the plastics material in the lower face 104 of a processing cartridge 43 as shown in Figures 5A-5C to aid correct alignment and to facilitate insertion of the processing cartridge during insertion of the cartridge into the shelf 41.

[0076] In the present example, where the guides 97 and 99 are formed from metal, the provision of the grooves 108 and 113 in plastics material at the upper and lower faces, respectively, of each information processing cartridge 43 (see Figures 5A-5B) results in a combination of metal and plastics materials that can provide a low friction interaction, facilitating insertion of the information processing cartridges.

[0077] If, for example, the information processing cartridge enclosure is made of a metal, it may be undesirable to provide metal guides to avoid a metal to metal interaction. In such a case, for example, it may be desirable to form the guides from a plastics material having a low coefficient of friction, such as polytetrafluoroethylene (PTFE) or polythene. Plastics rails could be attached to the underside of the plenum chamber floor member 94 and/or on the upper surface of the base 51 of the shelf 41. In such an example, grooves on the upper and lower faces of the information processing cartridges 43 could then be formed of metal or plastics and still result in a low friction arrangement.

[0078] A CSSP/PSU divider 96 can be provided to the rear of the midplane 171 and can extend horizontally to the rear face 59 of the shelf enclosure 47. The CSSPs 71, when inserted, are supported by the divider 96. To aid the correct insertion of the CSSPs 71, CSSP guide pins 178 are provided on the midplane 171 at positions adjacent connectors 177 on the midplane 171 for connection to the CSSPs 71.

[0079] Respective positions 88 and 89 can be formed in the front face 57 and the rear face 59 at which first and second indicator boards 183 and 184 supporting the indicator LEDs 69 can be located. These positions 88, 89 therefore include an aperture through the respective face of the shelf enclosure 47 such that indicator LEDs 69 mounted onto a circuit board attached to the inside of the shelf enclosure 47 may be viewed from outside the shelf enclosure.

[0080] There now follows are more detailed description of the midplane 171.

[0081] As mentioned above, the midplane 171 connects all the elements of a shelf together, including, in the present example, up to sixteen information processing cartridges 43, up to two CSSPs 71, two PSUs 81 and the two indicator boards 183 and 184. In the present example, due to its location within the shelf enclosure, the midplane 171 is not configured to be swappable. Accordingly, to maximize the system reliability, the midplane is configured to provide as a high level of reliability as possible. To this end, the midplane is advantageously configured without active devices and to include the minimum number of decoupling capacitors consistent with good design practice (ideally zero).

[0082] The midplane supports a number of paths for various power and signal lines to interconnect the FRUs.

[0083] In the present example, each information processing cartridge 43 has a high speed information signal connection (e.g., a Gigabit (Gb) Ethernet SERIALIZER/DESERIALIZER (SERDES) connection) to each of the CSSPs 71, each connection consisting of two pairs of differential signals. In a conventional manner therefore, the tracking of the paths for these signals is arranged to keep the pairs well balanced and on a single signal layer (i.e. without vias) to support such differential signals at high frequency.

[0084] In addition, in the present example, each information processing cartridge 43 has a serial console connection to the CSSP cartridge 71. Each connection consists of two TTL (Transistor-Transistor Logic) level signals that make a transmit and return (TX and RX) pair.

[0085] Also, each PSU 81 has a management signal connection (e.g., a serial I2C (Inter-IC Bus) connection) to the CSSP cartridge 71 to control power and monitor environmental parameters. The I2C bus comprises of two signals SCL and SDL (serial clock line and serial data line). In addition, an I2C address programming pin is provided for the PSUs 81.

[0086] Each information processing cartridge 43 and PSU 81 can signal to the CSSP cartridge 71 that it is inserted by pulling to ground (GND) a respective Inserted\_L signal (i.e., an active low signal). These signals are fed to the CSSP cartridge 71 via the midplane 171.

[0087] Each PSU 81 has five 12 Volt output rails. The routing from each PSU 81 is arranged so that a fault in any single FRU cannot completely interrupt the power to any other.

[0088] As mentioned above, the midplane 171 is provided with appropriate connector arrangements for receiving the connectors on the FRUs.

[0089] In the present example, the information processing cartridge 43 connects to the midplane 171 through a 40 pin Single Connector Attachment (SCA-2) connector as defined by the Small Computer Systems Interface (SCSI) standard. Accordingly, the midplane carries corresponding connectors 175.

[0090] In the present example, each CSSP cartridge 71 connects to the midplane 171 through a two right-angle 20 pair connector (e.g., 2mm HM-Zd connectors available from Tyco Electronics). The corresponding connectors 177 on the midplane are straight male parts with a power connector. A guide pin arrangement is provided in addition to the connectors to prevent misaligned modules causing bent pins during insertion. The guide pin also provides a leading ground. The CSSP cartridge 71 also connects to the midplane 171 through a right-angled 125 way 5 row 2mm connector. The connector 177 on the midplane 171 includes a straight male part. A guide pin arrangement is provided in addition to the connectors to prevent misaligned modules causing bent pins during insertion.

[0091] In the present example, as mentioned above, each PSU 81 connects to the midplane 171 through an SSI-MPS specification connector. The contacts are configured 5P/24S/6P with sequenced signal (S) and power (P) pins. Where the connector on the PSU is a 1450230-1 R/A male header, solder tails connector, the mating connector 179 on the midplane can be a 1450540-2 vertical receptacle, press-fit connector.

[0092] In the present implementation, indicator boards 183 and 184 (see Figure 8A) are provided at the front and rear of the system and are configured as FRUs. In this example they hold three system-level indicator LEDs 69 and include a FRU identity (FRU-ID) programmable read-only memory (PROM) each. Three LEDs 69 are present on the indicator board. There can, for example, be a white locator LED that can be switched by the user for locating the system; a green power-on LED to indicate when the system is powered; and an amber service-required LED to indicate a fault or other condition requiring servicing. These LEDs can be driven by the CSSP 71.

[0093] In the present example, identification information (FRU ID) for the midplane 171 is held on an I2C electrically erasable programmable read only memory (EEPROM) in the front indicator board 183. In addition to the I2C signals necessary to access the FRU ID EEPROM, the CSSPs 71 provide a current limited supply to the indicator boards 183 and 184 via the midplane. The indicator boards 183 and 184 are also provided with an I2C address programming pin. Depending on the implementation, FRU ID information can be stored instead, or in addition, on the rear indicator board 184.

[0094] As the FRU-ID for the midplane 171 is held on one or both of the indicator boards 183 and 184, the midplane can be a totally passive unit. The FRU-ID PROMs communicate with the CSSPs 71 via an I2C bus. Each device on the bus has a separate I2C address. The lower three I2C address bits of the EEPROMs used are available as pins on the device, to allow programming with resistors. The least significant bit of this address (A0) is passed to the midplane via the corresponding connector. This allows the midplane 171 to program the address of the FRU-ID differently for the front and rear indicator boards 183 and 184, by pulling the address low for the front board and high for the rear indicator board 183. This ensures that both EEPROMs are available on the bus, at different addresses. The FRU-ID for the midplane can be stored on either front or rear EEPROM, but the present example the FRU-ID is stored in the EEPROM on the front indicator board 183. The EEPROM can be 8kByte or larger.

[0095] As mentioned above, the midplane 171 includes openings 185 to provide a ventilation path for cooling air passing through the shelf 41. The cooling air passing through the shelf 41 via the midplane 171 can be driven by means of fans provided in each of the information processing cartridges 43 and the power supply modules 81. The openings 185 shown in Figures 8B,

9A, 9B and 9C form schematic representations of openings in the midplane 171. In practice, the openings could have any form (i.e., a series of large openings, or a number of small perforations), arranged on the midplane to align with corresponding openings or ventilation apertures in the various field replaceable units 43, 71 and 81. In this way, the path of the airflow from the front of the shelf to the back of the shelf can be configured to be as efficient as possible, depending on the detail configuration of the fan units and the ventilation openings or apertures in the information processing, switch, service processor and power supply unit modules 43, 71 and 81. Providing the fan units in the field replaceable units 43, 71 and 81, contributes to the aim of maintaining the chassis 49 and the midplane 171 of the shelf 41 free of active components, thereby minimising cost, and facilitating maintenance. Also, by providing the fan units in each of the field replaceable units, merely inserting and removing field replaceable units automatically adapts the flow of cooling air to the number and type of field replaceable units inserted in the shelf 41.

[0096] As described above, in the present example each of the FRUs is designed to be a non-user serviceable unit. Thus each FRU presents the user with a "sealed" unit which may be inserted into and removed from the shelf 41 as desired or required. If a FRU ceases to be operable, then the user has a choice only of returning the FRU to a supplier or service company for repair or of discarding the non-operable unit. As the FRUs are non-user serviceable, there is no requirement for a skilled technician to be employed in inserting or removing the FRUs into or from a shelf 41. Thus each FRU is designed such that a non-skilled person should have difficulty in causing damage to the FRU during handling. Moreover, the configuration and construction of the FRUs (e.g., the injector/ejector levers, the grooves in the enclosures of the information processing units, etc), of the shelf enclosure and the midplane (e.g., the guide rails to guide insertion of the FRUs, the locating pins, etc) contribute to facilitating easy insertion and removal of the FRUs.

[0097] Shown in Figure 10 is an example of the flow of cooling air through the shelf 41 and FRUs 43, 71, 81 mounted therein.

[0098] In this example, the cooling air passing through the shelf 41 is drawn generally in a front to rear direction through the shelf 41 by cooling fans mounted within the CSSPs 71 and the PSUs 81. Two separate flow paths for cooling air are provided in this example. The first, indicated as flow path  $\alpha\alpha$  by dotted lines 77 provides cooling air to the CSSPs 71. The second path, indicated as flow path  $\gamma\gamma$  by dotted lines 78 provides cooling air to the information processing cartridges 43 and PSUs 81.

[0099] The flow of cooling air along path  $\alpha\alpha$  enters the shelf 41 through the aperture(s) 65 in the front face 57 of the shelf enclosure 47 into the plenum chamber 66. Further apertures (not shown) could also be provided in the side faces 53 and 55 of the shelf enclosure 47 to

allow air to flow into the plenum chamber 66. This air then flows through the plenum chamber 66, and passes over the top edge of the midplane 171 to reach the perforations 133 of the front face of the CSSPs 71. The cooling air then passes through the CSSPs 71, providing cooling to the components thereof before passing out of the CSSPs 71 through the perforations 135 in the rear face of the CSSPs 71 thus being exhausted from the shelf 41. This flow of cooling air along flow path  $\alpha\alpha$  is driven by fans 79 mounted within the CSSPs 71. In the present example, a pair of fans 79 is provided within each CSSP 71 and is mounted against the rear face thereof.

[0100] Airflowing along path  $\alpha\alpha$  is impeded from flowing around the processing cartridges 43 by plenum chamber floor member 94 and is impeded from flowing to the PSUs 81 by CSSP/PSU divider 96. This flow path  $\alpha\alpha$  therefore ensures that air flowing to the CSSPs 71 is not warmed by passage through the processing cartridges 43 and therefore provides maximum efficiency cooling to the CSSPs 71.

[0101] The flow of cooling air along path  $\gamma\gamma$  enters the shelf 41 through the perforations 115 in the front face of the information processing cartridges 43. The air thus enters the information processing cartridges 43 and provides cooling to the components thereof. Cooling fans (not shown) within the information processing cartridges 43 direct the cooling air to the processor (CPU) of the information processing cartridge and direct the flow of air in the cartridge thereby increasing cooling efficiency. The air then exits the information processing cartridges 43 through the perforations 117 in the rear face thereof. The air then passes through the apertures 185 through the midplane 171 to reach the PSUs 81. This cooling air then passes through the perforations 157 on the front and upper front faces of the PSUs 81 to enter the PSUs and provide cooling to components thereof. It will be appreciated from Figure 10 that the sloping rear of the upper face of the PSUs 81 increases the area over which air can be drawn into the PSUs, thereby reducing the back pressure on the air flowing through the shelf unit and aiding the cooling efficiency. The flow of cooling air along path  $\gamma\gamma$  is driven by fans 85 mounted within the PSUs 81. In the present example, a pair of fans 85 is provided within each PSU 81 and are mounted against the rear face thereof.

[0102] Air reaching the PSUs 81 via path  $\gamma\gamma$  will already have passed through the processing cartridges 43. Such air will therefore be already warmed above the ambient temperature outside of the shelf 41 by its passage through the processing cartridges 43. However, as the cooling requirement of the PSUs 81 is typically less than that for the CSSPs 71, this does not cause any difficulty in the operation of the PSUs 81, which are adequately cooled by this flow of pre-warmed air. The pre-warmed air passing through the apertures 185 through the midplane 171 is impeded from flowing into path  $\alpha\alpha$  and entering the CSSPs 71 by the SCCP/PSU divider

96.

[0103] As will be appreciated by the skilled addressee, the arrangement shown in Figure 10 is illustrative only and other arrangements whereby, for example, a mixture of cool air from the plenum chamber 66 and warmed air from the processing cartridges 43 is provided to each rear-mounted FRU can readily be constructed.

[0104] With reference to Figure 11, there now follows a description of functional elements of an information processing cartridge 43 as contained within the information processing cartridge enclosure 101.

[0105] The information processing cartridge 43 includes a microprocessor 192 (a nonlimiting example of a microprocessor that can be utilised in the present example is an UltraSPARC™ processor). The microprocessor is mounted on an information processing cartridge motherboard 191.

[0106] A configurable clock generator 193, configured as a programmable clock synthesizer employing a crystal, can be used to produce CPU clock signals, CLKA and CLKB. The clock frequency can be determined by jumper settings (not shown). A vectored interrupt controller (I-Chip) 194 and a configurable core voltage regulator module (VRM) 195 are provided.

[0107] In the present example, memory means for use by the processor 192 when executing instructions can be provided in the form of buffered dynamic random access memory (DRAM), for example configured as dual in line memory modules (DIMMs) 196 with a 72-bit data path with error correction codes (ECC), seated in two sockets on a riser card from the information processing cartridge motherboard 191. The memory capacity can be chosen to suit the processor addressable memory space. For example, in the present example, up to 4 Gigabytes (4GB) of addressable memory can be provided. Serial Presence Detect (SPD) auto-configuration is provided via a Service Management Bus (SM-Bus) over an I2C bus 197.

[0108] In the present example, a PCI bus architecture can be employed with a so-called SouthBridge bus bridge 199 with SuperIO and two Gb Ethernet Media Access Control (MAC) devices. As described above, however, other bus protocols (e.g., Infiniband) can be used. A 32bit PCI bus 198 can be provided from the microprocessor 192. The SouthBridge 199 is a standard form of bus bridge, in the present example packaged in a 352 pin PBGA (Plastic Ball Grid Array) package, that provides the following functions: an SM Bus interface over the I2C bus 197 for access to the SPD (Serial Presence Detect) feature of the DIMMs that allows initialization of the memory controller; an Xbus interface for access via an Xbus 200 (which is a packet switched multiprocessor bus) to a PROM 201, a real time clock (RTC) 202 and an information processing cartridge service controller (hereinafter termed a Blade Service Controller (BSC)) 203; an IDE (Integrated Drive Electronics) interface that provides an ATA-100 (AT Attachment) IDE connection

204 to an IDE disk drive 205; and a serial console interface on a service bus 206 to the BSC 203 that is used for operating system functions including a console function with this embodiment.

[0109] For IO to the midplane 171, two AC-coupled Ethernet interfaces 207 and 208 are provided in the present example, which are packaged in a 316 pin PBGA. These Ethernet interfaces can provide a PCI attached Ethernet MAC capable of operation up to Gigabit Ethernet performance. The physical layer can be implemented using Serializer/Deserializers (SERDESs) 209 and 210. An example of a SERDES device is the TLK2201 transceiver manufactured by Texas Instruments, Inc. The SERDES devices use differential PECL TX+/- and RX+/- (Positive Emitter Coupled Logic Transmit and Receive) pairs to communicate to the switch portions of the CSSPs 71 over the midplane 171. The RX+/- pairs can be AC coupled at the information processing cartridge 43, the TX+/- pairs can be AC coupled at each CSSP 71. This facilitates hot-swap of the information processing cartridges 43 and the CSSPs 71.

[0110] Asynchronous serial connections 211 and 212 for communication between the BSC 203 and the Service Processor parts of the CSSPs 71 can be provided.

[0111] Internal data storage can be provided in the present example by a hard disk 205 with a capacity of 30GB or more rated for 24/7 continuous operation. The hard disk 205 is accessed using the primary IDE interface of the SouthBridge 199. The hard disk 205 can hold an operating system, for example a Solaris operating system, and other software and data for performing information processing using the main, or host, processor (CPU) within the information processing cartridge 43.

[0112] In the present implementation, the BSC 203 can be implemented as a microcontroller (e.g., a Hitachi H8 microcontroller). The BSC 203 can provide various functions, including for example: dual access (for the information processing cartridges and the CSSPs 71) to PROM 201 and EEPROM 213 for boot information and a FRU-ID for the information processing cartridge; channeling communication between an information processing cartridge 43 and the service processor part of the CSSPs 71; control of power on reset (POR), system reset and externally initiated reset (XIR) to the microprocessor 192; control of the power, service-required and ready-to-remove LEDs 69; upgrading of field-upgradable firmware, via the serial interface; a watchdog function for the operating system; monitoring the speed of a CPU fan 214; and communications with an EEPROM 215 and the operating system via the Xbus 200.

[0113] In the present example, the BSC 203 can be powered by a 5V service bus (SB) rail as soon as a CSSP 71 and a PSU 81 are fully inserted into the midplane 171, it then turns on other DC/DC converters to provide power to the remainder of the information processing cartridge 43. A BSC reset signal can be derived from a simple conventional power on reset (POR) generator that monitors a 5V supply rail.

[0114] In the present example a 1MByte Flash PROM 201 can be provided for storing boot variables for Open-Boot™ PROM (OBP) and Power-On-Self-Test (POST). Further OBP variables can be stored in a second 16kByte (16kB) I2C PROM 215, accessible via the SouthBridge SM Bus port over the IC Bus 197. The PROM 215 can contain 8kByte for OBP variables and 8kByte of unused space. A 16kByte I2C EEPROM 213 that is accessible via the BSC 203 can contain BSC variables and FRU-ID variables. The EEPROM is nominally divided into 8kByte for FRU-ID and 8kByte for the BSC variables. Write protection for the FRU-ID is implemented by BSC firmware. Such write protection may be carried out by, for example, acknowledging instructions to write to the protected area, but not to carry out those write instructions.

[0115] An environmental monitor sensor 215 can be provided to monitor the CPU and ambient temperatures. This sensor can be accessible via the onboard I2C bus from the BSC 203.

[0116] The information processing cartridge 43 can be powered from two, diode commoned, 9V power supply rails 216 and 217. DC/DC converters 218 can be used to provide the voltage levels required by the information processing cartridge 43. The DC/DC converters 218 are supplied by dual 9V inputs 216, 217, individually fused 219, 220 and then diode commoned 221, 222. A 5V DC/DC converter can be turned on as soon as the FRU is fully inserted, with the BSC 203 and required portions of the SouthBridge 199 being powered (the 5VSB rail). A field effect transistor (FET) can be used to gate off the main 5V supply to the rest of the information processing cartridge 43. The DC/DC converter outputs and the main 5V FET can be arranged not to turn on until the BSC 203 turns them on via a signal from the SouthBridge 199. The SouthBridge 199 can be used so that if the BSC 203 is reset (by a watchdog timeout or after a firmware download) the state of the DC/DC converters 218 is not affected. When the remaining outputs from the DC/DC converters 218 are within specification, a PWR\_GOOD signal can be asserted low to the BSC 203.

[0117] A SouthBridge resume circuit can be operable to run from 3V3, and a simple Zener diode dropper circuit can be used to generate 3V3 from the 5VSB supply.

[0118] When the FRU is inserted the inrush current can be limited, for example to <1A, and the rate of rise can be configured not to exceed a predetermined value (e.g., 20A/s) to provide a so-called soft start to facilitate hot-insertion. The intent is to prevent damage to the connectors and to avoid generating noise. A soft start controller 223, which controls a ramping-up of voltage levels, can be enabled when the predetermined signal (Inserted\_L signal) is asserted low, this signal is on a short pin in the connector and is connected to ground (GND - not shown) through the midplane 171.

[0119] In the present example, a processor impingement fan (processor fan) 214 is configured to run at full

speed to cool the information processing cartridge 43 and the fan. The speed of the processor fan and sink can be monitored by the BSC 203, using a tachometer sense pin on the microcontroller. In the event of the fan speed falling below a predetermined speed, or percentage of its nominal speed (e.g., 80%), the BSC 203 can be arranged to issue an alert. The nominal speed of the fan can be recorded as part of the BSC EEPROM contents.

[0120] The midplane connector 120 for the information processing cartridge 43 is used to establish the connection between the information processing cartridge 43 and the midplane. In the present example it supports up to 84 connections (pins) that will deliver SERDES outputs 224, 225, I2C signals 226, 227, and power 216, 217. Signal connections may be made through a right-angled connector. Power connections may be made through the information processing cartridge right-angled connector. The connector can be configured to facilitate hotswapping of the information processing cartridge, for example by having a low insertion force and/or guide pins to increase the ease of serviceability and prevent module misalignment during insertion.

[0121] Interrupts to the processor 192 can be encoded using an encoded interrupt vector mechanism. An I-Chip Emulator (ICE) 228 functions as an interrupt concentrator, receiving all system interrupts and encoding them as an interrupt vector according to an interrupt vector code utilisable by the processor 192. In the present example, where an UltraSPARC™ processor is used, the interrupt vector encoding may be based on a 6-bit interrupt vector code.

[0122] With reference to Figure 12, there now follows a description of an example of a combined switch and service processor (CSSP) 71. In the present example, each CSSP 71 provides the functionality of a Switch 73 and of a Shelf Service Processor, or Shelf Service Processor (SSP) 74.

[0123] Figure 12 provides an overview of the functional components of the CSSP 71 including functional components of the Switch 73 and functional components of the SSP 74. In the present example, most of the components relating to the Switch 73 are mounted on a Switch PCB 231, and the components relating to the SSP 75 are provided on a SSP PCB 232. However, it should be noted that the components located in the lower portion of the switch PCB 321 (i.e., that portion below the SSP PCB 232 as illustrated in Figure 12 logically belong to the SSP 74, rather than to the switch 73. It will be appreciated that such component arrangements are not compulsory for successful operation and that any other component arrangement over any number of component boards can be easily achieved using conventional component arrangement techniques.

[0124] Firstly, with reference to Figure 12, there follows a description of functional elements of the Switch portions 73 of a CSSP 71 as contained within the CSSP enclosure 121.



[0125] The midplane connector 141 on the CSSP 71 establishes the connection between the CSSP 71 and the midplane 171. In the present example, it supports up to 84 connections (pins) that will deliver SERDES outputs 265-268, I2C signals 310, 320, 321 and 322, and power 278, 279. Signal connections may be made through two 20-pair right-angled connectors. Power connections may be made through a right-angled connector. The connector can be configured to facilitate hot-swapping of the board, for example with a low insertion force. The connector also uses guide pins to increase the ease of serviceability and prevent module misalignment during insertion.

[0126] A switch microprocessor 240 is provided, in the present example the microprocessor used is a Power PC (MPC8245) packaged in a 352pin Tape Ball Grid Array (TBGA) package. This microprocessor 240 supports between 1MB and 2GB of address space in the present example. It further includes an Embedded Programmable Interrupt Controller (EPIC) that provides 5 hardware interrupts (IRQs) or 16 serial interrupts. There are 4 programmable timers with cascade mode function. DRAM memory for the processor can be provided in the present example by a commodity DIMM 242. The processor 240 can be connected to a 32bit PCI bus 241, which operates at, for example, 33MHz/66MHz.

[0127] A clock input to the processor 240 can be provided by a clock generator (CLK) 243. The CLK 243 can include a configurable clock generator (not shown) implemented as a programmable clock synthesiser employing a crystal used to produce CPU clock signals. The clock frequency can be determined by jumper settings (not shown). A vectored interrupt controller (I-Chip) (not shown) and a configurable core voltage regulator module (VRM) (not shown) can be provided that operate substantially as described above with reference to the like components of Figure 11.

[0128] In the present embodiment two switch ASICs (application specific integrated circuits) 244, 245 are provided (in the present example, BCM5632 Gigabit switch ASICs). Each ASIC can provide twelve GMII Interfaces (1 Gigabit Ethernet) (for uplinks and downlinks) and one 10Gb XGMII interface for chip-to-chip communication (bridging) 246 between the ASICs 244 and 245. Sixteen GMII 1Gb 'downlinks', in the form of serialized Gb Ethernet data, are provided through four quad SERDES 248-251 to allow each information processing cartridge 43 to communicate with the switch 73. Eight GMII 1Gb 'uplinks' are provided for external communication through two quad PHYs 253 and 254 (in the present example BCM5404 ASICs) and RJ45 connectors on the rear panel 122. The ASICs 244 and 245 are configured via a PCI interface (32bit/33MHz) to the PCI bus 241.

[0129] A Flash PROM 256 can store a real time operating system, and management and configuration data for the microprocessor. The Flash PROM 256 in the present example can be operable to hold 8MB - 16MB of data, depending on the software required. The flash

PROM 256 can be operated via an on-chip XBus 258.

[0130] Also connected to communicate with the processor 240 via the XBus 258, a Real Time Clock (RTC) 259 can be provided for real-time functions with a back-up battery.

[0131] Also connected to the XBus 258 can be a UART (Universal Asynchronous Receiver Transmitter) 260 which in turn connects to a serial bus 261 for providing an asynchronous console connection from the switch 73 to the SSP 74 which can be accessed by the SSP.

[0132] An integrated MAC/PHY (Media Access Control/Physical) switch 271 can provide its own interface to the PCI bus 241. This MAC/PHY switch 271 can connect to a 10/100 Ethernet hub 272. The hub 272 can be operable to provide a management interface to the SSP 74 and a connection from an external management network to the switch 73 and SSP 74 of a given CSSP 71. The connection from the integrated MAC/PHY device 271 to the SSP 74 can be coupled capacitively. A loopback mode can be provided by the MAC/PHY device 271 for system diagnostics. The hub 272 can connect to an RJ45 connector 273 on the rear panel 122 of the CSSP enclosure 121.

[0133] An 8kByte I2C EEPROM 262 can be used to store the FRU-ID and is accessible by the SSP portion 74 of each CSSP 71 via a serial bus 263 and the midplane 171. The upper 2kByte of the EEPROM 262 can be configured to be write protected.

[0134] An I2C Redundant Control Register (RCR) 275 can be used to provide an alternate, redundant path for powering-down the CSSP 71 and Shelf Level Indicators 69 mounted on the front 57 and rear 59 panels of the shelf 41. The I2C RCR 275 can be accessible by both the SSP 74 of the CSSP 71 containing the RCR and the SSP 74 of a further CSSP 71 connected via the midplane 171 via an I2C bus 276. In the present example, a device suitable for use as the RCR 275 is a Phillips PCF8574 IC.

[0135] With continued reference to Figure 12, there now follows a description of functional elements of the Shelf Service Processor (SSP) portion 74 of a CSSP 71 as contained within the CSSP enclosure 121 and provided on an SSP PCB 232.

[0136] In the present example, communication between the Switch PCB 231 and the SSP PCB 232 is facilitated by an interboard connector pair 298 and 299. It supports connections (pins) for I2C signals, 10/100 MAC/PHY output, and power. As described above, the switch PCB 231 carries the components associated with the switch, and it also carries the power, FRU-ID and environmental monitoring components along with the connectors for connections to the midplane 171 and external connectors. Thus, in the present example, all SSP components requiring a connection to the midplane 171 or an external connection have signal paths routed through the connector pair 298, 299 and via the switch PCB 231 to the relevant midplane or external connection.



tors.

[0137] In the present example, the SSP 74 includes a microprocessor 301 (e.g., a Power PC (MPC8245) processor) mounted on the SSP printed circuit board (PCB) 232. The processor 301 can be connected to a PCI bus 302, the present instance a 32 bit bus that operates, for example, at 33MHz/66MHz.

[0138] A clock input to the processor 301 can be provided by a clock generator (CLK) 303. The CLK 303 can comprise a configurable clock generator (not shown) implemented as a programmable clock synthesiser employing a crystal used to produce CPU clock signals. The clock frequency can be determined by jumper settings (not shown). A vectored interrupt controller (I-Chip) (not shown) and a configurable core voltage regulator module (VRM) (not shown) can be provided that operate substantially as described above with reference to the like components of Figure 11.

[0139] The processor 301 can be provided with a DRAM memory 305. The memory capacity can be chosen to suit the processor addressable memory space. In the present example, 8 MB of DRAM memory is provided.

[0140] An integrated MAC/PHY switch 306 can provide its own interface to the PCI bus 302. The MAC/PHY switch 271 can be connected to 10/100 Ethernet hub 272 via the interboard connectors 298, 299. A loopback mode can be provided by the MAC/PHY switch 306 for system diagnostics.

[0141] Octal UARTs 308 and 309 can be connected between the PCI bus 302 and the interboard connector pair 298, 299. The signal path can be continued from the interboard connector pair 298, 299 to serial connections 310 on the midplane connector 141 on switch PCB 231. The Octal UARTS 308, 309 can facilitate serial communications between the SSP 74 and each of the processing cartridges 43.

[0142] Also connected to the PCI Bus 302 can be a dual UART (DUART) 312 that in turn can connect via the interboard connectors 298, 299 to serial bus 261 for providing an asynchronous console connection from the SSP 74 to the switch 73. The DUART 312 can also have an I2C connection to an external connector on the rear face 122 of the CSSP enclosure 121. The external connector can provide a common operating system/boot console and command port 311.

[0143] Connected to the processor 301 via an XBus 314 can be a Flash PROM 315. The Flash PROM 315 can store a real time operating system, and management and configuration data for the microprocessor 301. The Flash PROM 315 can be operable in the present example to hold up to 2MB of data, depending on the software required.

[0144] Also connected to the processor 301 via the XBus 214 can be a real time clock (RTC) 316 for real-time functions with a backup battery. The RTC 316 can also provide 8kByte of non-volatile random access memory (NVRAM), in the present instance implemented

as an EEPROM. This can be used to contain information such as the FRU-ID, a serial number and other FRU information.

[0145] To facilitate I2C communications between the SSP 74 and the other CSSP 71, the midplane 171 and the PSUs 81, a multiplexer 318 can be provided. The multiplexer 318 can have a single I2C connection to the processor 301 and connections, via the interboard connector pair 298, 299 and the midplane connector 141 to both PSUs 81, the midplane 171 and the other CSSP 71.

[0146] The processor 301 can also comprise an embedded DUART to provide a redundant serial link to the SSP 74 of the other CSSP 71. Although it would be possible to implement this link using an external DUART, the advantage of using an embedded DUART is that the connection to the other CSSP is reliable and therefore likely to be functional. Where the embedded DUART link does not use the I2C Multiplexer for communications to the other CSSP, a common mode of failure for both the SSP - SSP I2C links can be avoided, it being assumed that the processor 301 is likely to be functional even if both embedded DUART channels are non-functional.

[0147] The CSSP 71 can be powered from two, diode commoned, 9V power supply rails 278 and 279. DC/DC converters 281 can be used to provide the voltage levels required by the CSSP 71. The DC/DC converters 281 can be supplied by dual 9V inputs 278, 279, individually fused 285, 286 and then diode commoned 287, 288. A soft start controller 283 can be provided to facilitate hot-insertion. A 5V DC/DC converter (I2C power regulator) 282 can be turned on as soon as the CSSP 71 is fully inserted. A 3.3V DC/DC converter can be turned on when instructed, for example through SSP service software, by asserting low an appropriate signal (ON\_L - not shown). The 3.3V converter can be arranged to turn on a converted for 2.5V, 1.2V, and a processor core voltage rail (Vcore) when the voltages are within an appropriate range.

[0148] When the CSSP 71 is inserted the inrush current can be limited, for example to <1A, and the rate of rise can be configured not to exceed a predetermined value (e.g., 20A/s) to provide a so-called soft start to facilitate hot-insertion. The intent is to prevent damage to the connectors and to avoid generating noise. A soft start controller 283, which controls a ramping-up of voltage levels, can be enabled when the predetermined signal (Inserted\_L signal) is asserted low, this signal is on a short pin in the connector and is connected to ground (GND - not shown) through the midplane 171 until one of the supplies is removed. These circuits can be configured to withstand an overvoltage at their inputs whilst the input they are feeding is not powered, without any leakage to the unpowered circuit. A sense circuit can detect if the voltage has dropped below a threshold, for example 2.0V, as a result of a blown fuse, a power rail going down, etc. The DC/DC converters 281 can be protected against short circuit of their outputs so that no damage occurs.

[0149] The I2C regulator 282 can be powered as soon as the CSSP 71 is fully inserted into the midplane 171. This can be facilitated through short pins connected to the soft start controller 283, which controls a ramping-up of voltage levels. The other DC/DC regulators can be turned on, for example by SSP software.

[0150] A pair of fans 290, 291 can provide cooling to the CSSP 71. The fans 290, 291 can be configured to run at full speed to prevent overtemperature conditions by minimizing the temperature of the internal components and the fan. The speed of the fans 290, 291 can be monitored by the SSP 74 through an environmental monitor 295 on the switch board 231. The environmental monitor 295 can be alerted in the event of the fan speed falling below a predetermined value (e.g., 80% of its nominal speed). The fan can provide tachometer outputs to facilitate the measurement of fan speed.

[0151] LED indicators 137 can be provided, for example with a green power LED, an amber LED for indicating that service is required and a blue LED for indicating that the switch is ready to be removed. LED indicators integrated on 2x4 stacked RJ45 connectors on the rear face of the CSSP 71 can be arranged, for example, to show green continually when the link is present and flash green when the link is active.

[0152] The environmental monitor ENV MON 295 can be provided to maintain operational integrity of the CSSP 71. The ENV MON 295 can include limit values in limit registers and can monitor, for example, temperature within the CSSP enclosure 121, the CSSP power rails, including the 12V, 3V3, Switch Processor Core Voltage, CSSP Processor Core Voltage and the two 9V power feed rails 278, 279 from the midplane 171. The outputs of the DC/DC converters 281 can be fed in to A/D inputs of the ENV MON 295 for Watchdog comparisons to be made to the voltage limits set in the limit registers. As noted above, the ENV MON 295 can also monitor the operating speeds of the fans 290 and 291. The ENV MON 295 can communicate with the SSP 74 of both CSSPs via an I2C bus 296.

[0153] For IO to the midplane 171 shown in Figures 8A-10, the midplane connector 141 can include sixteen 1Gb Ethernet connections 265-268 from four quad SERDES 248-251 and the I2C bus lines 596.

[0154] The SSP 74 can access the I2C devices (FRU-ID EEPROM, 8-bit I/O expansion chip, and the system hardware monitor) through the midplane 171.

[0155] For external IO, rear panel Gb Ethernet connections can be provided from the two quad PHYs 253, 254 to 2x4 stacked RJ45 connectors 139 (to give 8 uplinks). Each port can be an independent 10/100/1000 BASE-T (auto negotiating) port. The PHY devices 253, 254 can operate in GMII mode to receive signals from the 8-Gigabit interfaces on the ASICs 244, 245.

[0156] The Power Supply Units (PSUs) 81 can be configured such that when two or more PSUs 81 are connected in parallel in the shelf 41, failure of any one of the paralleled units shall not affect system operation.

Moreover, one of the PSUs can be installed or removed from a "live" system with or without input power applied. The outputs can have overcurrent protection.

[0157] The PSU can have an I2C interface to provide power supply status via the midplane 171. The PSU can have an internal temperature sensor that reports via the I2C interface. The PSU fan speed can also be monitored and errors are reported via the I2C interface. Overvoltage and overcurrent sensors can also report via the I2C interface.

[0158] There now follows a description of aspects of an example of a power supply 81 with particular reference to Figure 13.

[0159] When a power supply (e.g. mains power, or UPS type protected power) is connected to the cable connector 83, transformer, regulator and rectifier circuitry 400 can operate to generate a DC output (in the present example, 12V DC) from the input (in the present example 230/240V 50Hz AC or 110V 60Hz AC).

[0160] In order to ensure good cooling reliability within the shelf 41, each PSU 81 can have a pair of cooling fans 402, 403 located at the rear of the PSU enclosure as described above with reference to Figure 7. As a fail-safe measure against a failure of one PSU 81, the fans of each PSU 81 can be powered by both PSUs 81. Thus, in the event that one PSU 81 is non-operative to produce the DC supply for the shelf 41, as long as the other PSU 81 is operative, not only will all components of the shelf 41 be powered as normal, the fans of both PSUs 81 can continue to run.

[0161] As shown in Figure 13, this dual powering of cooling fans 402, 403 can be effected by providing a power supply line 404 from the transformer, regulator and rectifier circuitry 400 to power both fans 402, 403. Also, first and second separate power lines 410, 412 from the other PSU 81 can provide duplicate power supply to the first and second fans 402, 403 respectively. The fan 402 can thus be powered by a diode commoned supply from line 404 and a diode commoned supply from line 410. Diode protection can be provided by diodes 405 and 411 respectively. The speed of the fan 402 can be controlled by a speed controller 408. Similarly the fan 403 can be powered by a diode commoned supply from line 404 and a diode commoned supply from line 412. Diode protection can be provided by diodes 406 and 414 respectively. The speed of the fan 403 can be controlled by a speed controller 409.

[0162] The two speed controllers 408, 409 can in turn be controlled by a data input from each CSSP 71 received via an I2C bus connection (not shown in Figure 13). Power supply lines carrying DC power for the other FRUs of the shelf 41 are shown in Figure 13 as power line 416. All power connections to and from the PSU 81 can connect to the midplane 171 when the PSU is inserted in the shelf 41 via the midplane connector 163. In the present example the PSU 81 connects to the shelf through a 5P/24S/6P configuration SSI-MPS compliant right angle connector 163 at the front face 147 of the

PSU 81. Connectors for the I2C interface can also be provided.

[0163] The input power line 410 and 412 for each fan 402 and 403 can be provided with a softstart module 4131 and 4132 respectively, to allow for hot insertion of the PSU 81 into the shelf 41. The softstart modules 4131 and 4132 can be controlled, for example, by pulling a signal to ground (e.g., a "mated" input line 4151 and 4152).

[0164] Where the two input power lines 410 and 412 are separate lines having separate softstart provision, there is no common failure mode for the backup method of powering the fans 402, 403. Thus even if the PSU 81 ceases to be operable to generate the DC supply, and a component (power line or softstart module for example) fails in the supply path from the second PSU 81 to the fans 402, 403, at least one of the fans 402, 403 can remain operational as the shelf 41 still receives the cooling effect of three PSU fans.

[0165] In the present example, the power supply has four rear panel LED indicators 137. A blue "Ready to Remove" LED can be driven by the I2C interface and indicate that the power supply may be removed from the system. An amber "Service Required" LED can be driven by the I2C interface and indicate that the power supply is in a fault condition: any output out of range, over-temperature or shutdown. A green "DC Output-OK" indicator can be driven by internal power supply circuits and show that the main 12 volt supply is functioning. The LEDs can remain lighted when individual outputs are in the current limited mode of operation. A green "AC Input-OK" indicator can be driven by internal power supply circuits and show that AC input power is within normal operating range.

[0166] With reference to Figure 14, there will now be described an example of data connectivity between the FRUs and midplane of the shelf 41. Power transmission paths are not illustrated in Figure 14. However, it will be appreciated that to facilitate the maximum component redundancy of the shelf 41, each PSU 81 can independently provide power to each FRU.

[0167] In the present example each of the processing cartridges (blades) 43 connects to the midplane 171 via a pair of information signal connections (e.g. Gb Ethernet links) 224, 225 and a pair of serial management signal connections 226, 227. Connections within the midplane 171 can ensure that each Ethernet link 224 is directed to a connection 265-268 from the midplane 171 to a first switch 73, and that each Ethernet link 225 is directed to a connection 265-268 from the midplane 171 to a second switch 73. Thus one Ethernet link can be established between each processing cartridge 43 and the switch 73 of each CSSP 71. Further connections within the midplane 171 can ensure that each serial connection 226 is directed to a connection 310 from the midplane 171 to the first SSP 74 and that each serial connection 227 is directed to the second SSP 74. Thus one serial link can be established between each processing

cartridge 43 and the SSP 74 of each CSSP 71. As mentioned earlier, information signal connections other than Gb Ethernet connections (e.g. Infiniband connections) could be employed in other examples.

[0168] A plurality of serial connections can connect each SSP 74 to the other. Serial lines 320, 321 can connect each SSP 74 to the midplane 171 and connections within the midplane can connect the two sets of lines together. To provide a control interface from the SSPs 74 to the PSUs 81, serial lines 322 can connect each SSP 74 to the midplane 171 and connections within the midplane 171 can connect to serial lines 324 from the midplane 171 to each PSU 81.

[0169] An example of data and control connectivity of the shelf 41 to and from computer systems external to the shelf 41 when the shelf 41 is arranged for use within a multiprocessor server system such as that described above with reference to Figure 1 will be described with reference to Figure 15.

[0170] As summarised above with reference to Figure 14, in the present example each processing cartridge, or blade, 43 is connected to the switch 73 of each CSSP 71 by an information signal connection (e.g. a 1Gb Ethernet link) formed by a combination of links 224, 225 from the processing cartridge 43 to the midplane 171, connections within the midplane 171 and links 265-268 from the midplane 171 to each switch 73.

[0171] Further, in this example a set of serial management signal connections comprising links 320, 321 and connections within the midplane 171 connect the SSP 74 of each CSSP 71 to the SSP 74 of the other CSSP 71.

[0172] To provide external data connectivity between the shelf 41 and an external core data network 330, in association with which all information processing performed by the processing cartridges 43 of the shelf 41 is undertaken, connections 331 can be formed between the core data network 330 and the eight 1Gb Ethernet ports 139 provided on the rear panel 122 of the CSSP enclosure 121.

[0173] In the present example, the connections by means of which control and configuration of the shelf 41 are performed are entirely separate to the connections to the core data network 330. Therefore, a first external switch 335 can connect to a management (I2C) port 273 of the first CSSP 71 and a second external switch 336 can connect to a management (I2C) port 273 of the second CSSP 72. As described above with reference to Figure 12, the management port 273 can provide a management network interface to both the switch 73 and SSP 74 of each CSSP 71. The external switches 335, 336 can each be connected to each of a pair of System Management Server (SMSs) 338, 339. The SMS is not essential to the operation of the shelf 41, but use thereof aids optimal operation of the shelf 41. In a typical multiprocessor server system a plurality of shelves 41 may be connected together via the core data network 330 under the control of a single management network utilising one set of SMSs 338, 339. A set of SMSs 338,

339 may comprise a single SMS (as well as a plurality thereof). However use of at least two SMSs enables redundancy of components, therefore increasing overall system reliability.

[0174] A serial interface control 343 operable under telnet protocol control is also connected to the shelf 41 in the present example. This can provide a common operating system/boot console connection to the SSP 74 of both CSSPs 71 via the RJ45 connector 311 on the rear panel 122 of each CSSP enclosure 121.

[0175] It will be appreciated from the above that a flexible and scalable modular computer architecture has been described. In the described example up to 16 information processing cartridges, or blades 43, can be configured as sealed FRUs on a single shelf 41, the number of blades being chosen according to customer requirements. Each blade has its own processor and random access memory. If, for example, there is a maximum of 2Gbytes of memory per information processing cartridge, and one processor per blade, 16 processors (16P) with 5.33 processors per unit height (1U) and a total of 32GB of memory per shelf can be provided.

[0176] In the present example, the shelf 41 incorporates redundant combined switch and shelf service processor modules (CSSPs) 71 and redundant power supply units (PSUs) 81 separate from the blades 43. As the power supplies are carried by the shelf, the information processing cartridges can be kept compact and inexpensive. Also, as a result, they can be powered by DC power only, via the midplane 171.

[0177] Also, as mentioned earlier, the FRUs (e.g., the information processing cartridges, or blades, 43, the CSSPs 71 and the PSUs 81) can all be configured as sealed units that do not contain any internal FRUs themselves and do not contain user serviceable items. The enclosures of the FRUs can be arranged to enclose all of the functional components of the FRU with only electrical connectors being externally accessible and with indicator LEDs being externally visible as well.

[0178] These factors can all contribute to keeping the cost of the FRUs low, as well as that of the overall system. The modular approach with the use of sealed modular field replaceable units for providing system functionality and with non-field replaceable units designed with a minimum possible number of active components enhances reliability. Moreover, easy and rapid maintenance is facilitated in the event of a failure of a FRU by simple replacement of that FRU, further reducing the cost of ownership.

[0179] Thus, it will be appreciated from the above description that the provision of a rack mountable shelf, that includes power supplies, a shelf service processor and switches in modular units, for carrying a number of processing cartridges, wherein the number of processing cartridges can be chosen according to customer requirements, provides a flexible and scalable computer configuration. The balancing of the load between the processors of the processing cartridges can be effected

by software using conventional principles.

[0180] A configuration as described provides an easily scalable processor architecture, whereby the processing power provided by a complete system based on the information processing cartridge/information processing cartridge carrier architecture can be scalable from moderate to very high capacity through the simple addition of further information processing cartridges.

[0181] As illustrated in Figure 16, an example of the external connections from a shelf 41 can be in the form of two active information signal connections (e.g., Ethernet connections) 350 and 351, two active power connections 353 and an active/standby pair of management connections 354. With regard to the management connections, each connection comprises a serial connection and a network (e.g. Ethernet or Infiniband) connection. It is possible to connect to either the active or the standby connection, as the incoming signal will be internally routed to whichever management controller (CSSP) is the current master. It will be appreciated, therefore, that the connections to a shelf can be kept to a minimum. It will further be appreciated from the configuration shown in Figure 16 that the system is scalable beyond a single shelf unit 41.

[0182] Figure 17 illustrates how a plurality of shelves can be configured within one (or more) racks to provide even higher processing power. Such a constellation of shelves to provide a large grouping of servers is sometimes termed a "web farm" or "server farm" 360. As shown in Figure 17, the web farm comprises a plurality of shelves 41 that each carry a plurality of blades 43. Also provided are a plurality of Network Attached Storage devices (NAS) 373 for providing storage for critical data, e.g. email data storage, for the web farm. The NASs 373 are not required if there is no critical data to be stored, e.g. if the web farm is operating solely to provide web caching services.

[0183] Management control of the web farm 360 can be provided through a pair of System Management Servers (SMSs) 362. Each SMS 362 can be connected to a management network via a link 366 and to a management console 365. The SMSs 362 can communicate with the individual shelves 41 via a pair of management switches 364. Each shelf 41 and NAS 373 can be connected to each management switch 364 via a connection 367. Thus dual redundant management connections can be provided to each shelf 41 and NAS 373.

[0184] Flow of data to and from the web farm 360 can be provided through a pair of data switches 369. Each data switch 369 can be connected to a consumer network via a link 370. It is to be understood that the consumer network can be a larger data network to which the web farm 360 is connected. This network can be an office or corporation intranet, a local area network (LAN), a wide area network (WAN), the Internet or any other network. Connections between the data switches and the shelves 41 can be facilitated by connections 371. It is to be noted that as each shelf has its own

switching capability, there is no need for each shelf 41 to be directly connected to the data switches 369. Connections can also be provided to connect the NAS units 373 to the shelves 41. The topology used for interconnection of the data switches 369, shelves 41 and NASs 373 can be any topology providing at least one connection of any length between every possible pair of units. Complex topologies arranged to minimise the maximum connection length between any two given units in the web farm can be used.

[0185] The web farm 360 comprising a plurality of shelves 41 with or without a plurality of NASs 373 can suitably be used as any or all of the entry edge server group 9, web edge server group 15 and application servers 19 described above with reference to Figure 1.

[0186] As an alternative to providing critical data storage within a NAS 373, such storage can be provided within one or more NAS cartridges fitted into one or more of the shelves 41 in place of processing cartridges 43. Another alternative is to provide a server shelf with local storage such as a RAID array (Redundant Array of Inexpensive Disks) in place of the NAS 373.

[0187] Thus there has now been described an example of a fully configurable computing system based on a plurality of self contained field replaceable units (FRUs) and scalable from a single processing cartridge with power supply and switching capability to a multiply redundant multiprocessor server system with full system management capability extending over a number of co-operably connected server shelves. It will of course be readily apparent to the skilled reader that many of the specific features specified in the above description are in no way limiting and a variety of alternatives may be produced using only ordinary skill and common general knowledge. Non-limiting examples of example modifications which may be made to the above described system are discussed hereafter.

[0188] There is no limit placed on the processing cartridges as to what software they should run. Each module within a shelf or farm may run under the same operating system, or a plurality of different operating systems may be used. Examples of possible operating systems include Sun Microsystems' Solaris® OS or another UNIX™-Type OS such as Linux™, MINIX™, or Irix™, or UNIX™ or a Microsoft OS such as Windows NT™, Windows 2000™, Windows ME/98/95™, Windows XP™.

[0189] It is also not necessary that each processing cartridge within a shelf or farm be configured to run the same program software. For example, individual processing cartridges may be configured to execute, for example, fileserver software, mailserver software, web-hosting software, database software, firewall software, or verification software.

[0190] Although it has been described above with reference to Figures 4, 8, 12, 13 and 14, that functionality of a switch and of a shelf service processor should be provided within a single combined switch and service processor unit, this is not essential and separate switch

and shelf service processor field replaceable units may be used.

[0191] Although it has been described above that a pair of PSUs and a pair of CSSPs may be provided so as to enable dual-redundancy, further PSUs and CSSPs may be provided so as to increase FRU redundancy further, thus providing statistically higher reliability.

[0192] In the power supply circuitry in each of the blades and CSSPs, two voltage sense circuits may be provided after the fuses and before the diodes, to prevent a latent fault caused by a failed fuse going undetected until one of the PSUs is removed or taken offline. Such circuits may be configured to withstand an overvoltage at their inputs whilst the input they are feeding is not powered, without any leakage to the unpowered circuit.

[0193] Although it has been described above with particular reference to Figure 11 that the processing module may be based on an UltraSPARC™ processor, this is not limiting and any other processor having sufficient processing capacity to undertake the tasks required of a particular processing cartridge may be used. Alternative processors include, but are not limited to, Intel x86 series and compatible processors, AMD x86 compatible processors, Alpha processors and PowerPC processors. The particular example of an x86 compatible processor is described in more detail with reference to Figure 18. In Figure 18, the parts corresponding to those of the UltraSPARC™ based system of Figure 11 have the same reference numerals and will not be described again here. In a system based on an x86 compatible processor, the processor 378 itself communicates with the other components, including the memory 196 and PCI bus 198 via a Northbridge 379. The Northbridge 379 also includes an interrupt controller, so no separate interrupt concentrator is required. The other components of the processing cartridge could be substantially the same as for the UltraSPARC™ based system described above.

[0194] Although it has been described above that each information processing cartridge comprises a single microprocessor, this is not a limiting case as each or any of the information processing cartridges may have more than one microprocessor arranged to share common storage resources to operate synchronously (in lockstep) or asynchronously. Also, it is not necessary that all information processing cartridges inserted into a shelf at a given time are identical, rather a variety of different blade architectures may be used simultaneously.

[0195] The provision of the functions of both Switch and Shelf Service Processor within a single FRU in the present example provides a facility within a single shelf 41 for dual redundancy in both functions in fewer different FRUs. As will be appreciated, there is no restriction that these two functions must be provided within a single FRU and division of the two functions into separate FRUs would present no difficulty to the skilled addressee.

[0196] The backup power provision for the cooling fans of the PSUs is described above with reference to Figure 13. Although it is described that the backup power supply to each fan should be independent of the other, if the risk of common mode failure of backup power supply is judged to be low, or is of low importance, then the backup supply to each fan may be provided in common with the backup supply to all other fans.

[0197] Also, an information processing module forming a field replaceable server blade can include a processor and memory can be configured by means of software, firmware or hardware to provide a special purpose function. By way of examples only, an information processing module can be configured to perform the function of one or more of a firewall, or a load balancer, encryption and/or decryption processing, an interface to a secure network, e.g. a virtual private network (VPN), a specialized switch with wide area network (WAN) connectability.

[0198] Also, a storage blade may be provided. The storage blade can be configured to be mountable in a server blade receiving location in a blade server carrier. The storage blade can comprise storage blade connectors configured for interconnecting with carrier connectors on the server blade carrier, whereby the storage blade is interchangeable with a server blade. A carrier, or shelf, for a server system can be arranged with a plurality of blade receiving locations for receiving blades, wherein the blades can be storage blades or information processing blades. The server system can be self configuring on receipt of the blades according to the type of blade received in each said location. To achieve the blade service controller in each blade can be operable to communicate with a shelf service processor to perform said configuring.

[0199] In the above described systems, it is proposed to provide two information connections from each information processing module, or server blade, 43 over the midplane to the respective CSSP modules 71. The provision of these redundant information connections means that in the event of the failure or absence of one of the CSSPs 71, the information processing module 43 can continue to function using the other of the information connections. It will be appreciated that more than two information connections could be provided to yet further increase levels of redundancy. However, although such additional information connections might be provided in exceptional applications, the cost of such additional information connections would not be justified in most applications as the additional reliability gains would typically be very small compared to the additional cost.

[0200] In the examples of an information processing module described with reference to Figure 11 or Figure 18, information signal interfaces that include an Ethernet interface and a serializer/deserializer are illustrated. Figure 19A reproduces part of the information processing module shown in Figure 11. Figure 19A shows an

example where, for enabling information to be passed between the processor 192 to the midplane 171, two information protocol interfaces (IPIs) in the form of Ethernet protocol interfaces (EI) 207 and 208 are provided. The Ethernet protocol interfaces 207 and 208 are connected to the processor 192 via the processor bus 198 and respective bus connections 411 and 412. Each of the Ethernet protocol interfaces 207 and 208 is also connected to a respective physical layer interface (PLI) 209 and 210 in the form of SERIALIZER/DESERIALIZER (SERDES) devices by respective information signal paths 413 and 414. The SERDES devices 209 and 210 are connected in turn to respective information connections 224 and 225 on the midplane connector 120 via respective information connection paths 415 and 416. The combination of the IPIs (Ethernet protocol interfaces 207, 208) and the PLIs (SERDES devices 209, 210) forms an information signal interface 401.

[0201] As has already been described above, the use of the SERDES devices 209 and 210 means that the PLI is able to support not only an Ethernet information protocol, but also an Infiniband information protocol. In this case, the information connection paths 415 and 416 will each comprise differential PECL TX+/- and RX+/- (Positive Emitter Coupled Logic Transmit and Receive) pairs to communicate to the switch portions of the CSSPs 71 over the midplane 171. The RX+/- pairs can be AC coupled at the information processing module 43 and the TX+/- pairs can be AC coupled at each CSSP 71. The SERDES devices are able to support multiple information protocols, and in particular the higher bandwidth requirements of an Infiniband protocol compared to an Ethernet protocol. Accordingly, in order to adapt the information processing module for operating under an Infiniband protocol it is merely necessary to replace the Ethernet protocol interfaces 207 and 208 by Infiniband protocol interfaces. This is illustrated in Figure 19B.

[0202] Figure 19B corresponds to Figure 19A with the exception that the Ethernet protocol interfaces 207 and 208 of Figure 19A have been replaced by Infiniband protocol interfaces (II) 407 and 408. The Infiniband protocol interfaces 407 and 408 are connected to the processor 192 via the processor bus 198 and the respective bus connections 411 and 412. Each of the Infiniband protocol interfaces 407 and 408 is also connected to a respective one of the physical layer interfaces (PLI) 209 and 210 in the form of SERIALIZER/DESERIALIZER (SERDES) devices by the respective information signal paths 413 and 414. The SERDES devices 209 and 210 are connected in turn to respective information connections 224 and 225 on the midplane connector 120 via respective information connection paths 415 and 416.

[0203] Accordingly, the use of SERDES devices as physical layer interfaces enables the information processing modules to be supplied with a minimum of modification to support an Ethernet protocol, Infiniband protocol, or indeed any other appropriate information

protocol supportable by the SERDES devices.

[0204] The examples of Figures 19A and 19B illustrate the different configurations of the information signal interfaces for an information processing module as shown in Figures 11 or 18 that would typically be effected prior to delivery to a customer to provide support an Ethernet protocol or an Infiniband protocol, respectively. However, it may also be desirable to provide a generic information processing module that could be adapted at initialization or in use for an appropriate information protocol. Figures 20A and 20B illustrate alternative configurations that can enable configuration of a generic information processing module at initialization or in use for an appropriate information protocol. Figure 20A and 20B illustrate the use of Ethernet and Infiniband protocol interfaces although it will be appreciated that other examples of the invention could be provided with interfaces for other information protocols as appropriate for a given application.

[0205] It will be noted that in the examples of Figure 20A and 20B, an Ethernet protocol interface 207, 208 and an Infiniband interface 407, 408 are provided for each information connection 224 and 225 on the midplane connector 120, with switching logic, for example one or more multiplexer/demultiplexers, to enable the selection of the appropriate information protocol interface.

[0206] In the example shown in Figure 20A, the processor 192 is connected via respective signal paths 483 and 484 to first and second IPIs 207 and 407, respectively. The first IPI can be, in this example, an Ethernet protocol interface (EI) 207 and the second IPI can be, in this example, an Infiniband protocol interface (II) 407. The Ethernet protocol interface 207 can be connected via an information signal path 467 to a PLI 209, which can be implemented by a SERDES device as described earlier. The Infiniband protocol interface 407 can be connected via an information signal path 468 to a PLI 409, which can be also be implemented by a SERDES device. The SERDES devices 209 and 409 can each be connected via respective signal paths 487 and 488 (i.e., respective PECL TX+/- and RX+/- signal line pairs from the SERDES devices 209 and 409) to respective inputs of switching logic (SL) 440 (in the present example a 4:2/2:4 multiplexer/demultiplexer providing 2:1 multiplexing and 1:2 demultiplexing for each signal line of the differential PECL TX+/- and RX+/- signal line pair output by the SERDES devices 209 and 409). The switching logic 440 thus provides a common information signal path 472 (i.e. respective PECL TX+/- and RX+/- signal line pairs) to an information connection 224 on the midplane connector 120. The switching logic 440 can be controlled via a common switching signal line (not shown) from, for example, the information processing module service controller (BSC 203) shown in Figure 11. The combination of the IPIs (e.g., Ethernet protocol interface 207 and Infiniband protocol interface 407), the PLIs (e.g., SERDES devices 209 and 409), and the

switching logic (e.g., multiplexer/demultiplexer 440), forms an information signal interface 431.

[0207] Further, in the example shown in Figure 20A, the processor 192 is also connected via respective signal paths 485 and 486 to third and fourth IPIs 208 and 408, respectively. The third IPI can be, in this example, an Ethernet protocol interface (EI) 208 and the fourth IPI can be, in this example, an Infiniband protocol interface (II) 408. The Ethernet protocol interface 208 can be connected via an information signal path 469 to a PLI 210, which can be implemented by a SERDES device as described earlier. The Infiniband protocol interface 408 can be connected via an information signal path 470 to a PLI 410, which can be also be implemented by a SERDES device. The SERDES devices 210 and 410 can each be connected via respective signal paths 489 and 490 (i.e., respective PECL TX+/- and RX+/- signal line pairs from the SERDES devices 210 and 410) to respective inputs of switching logic 442 (in the present example a 4:2/2:4 multiplexer/demultiplexer providing 2:1 multiplexing and 1:2 demultiplexing for each signal line of the differential PECL TX+/- and RX+/- signal line pair output by the SERDES devices 210 and 410). The switching logic 442 thus provides a common information signal path 473 (i.e. respective PECL TX+/- and RX+/- signal line pairs) to an information connection 225 on the midplane connector 120. The switching logic 442 can also be controlled via a common switching signal line (not shown) from, for example, the information processing module service controller (BSC 203) shown in Figure 11. The combination of the IPIs (e.g., Ethernet protocol interface 208 and Infiniband protocol interface 408), the PLIs (e.g., SERDES devices 210 and 410), and the switching logic (e.g., multiplexer/demultiplexer 442), forms an information signal interface 432.

[0208] On initialization, or during use, the information processing module service controller 203 can be operable to select one of the information protocol interfaces by sending a switching signal to the switching logic 440 and 442. The signal can be generated in response to user input at a management console connected through the hierarchical management structure described earlier, or in response to an automatic exchange of information between the service controller 203 in the information processing module 43 and the shelf service processor 74 in the CSSP 71.

[0209] For example, initialization software or firmware in the information processing module, for example in the PROM 201 shown in Figure 11, can be operable on initialization to cause the service controller 203 to send management information to the shelf service processor 47 to inform the shelf service processor 74 that it has a configurable information signal interface able to support given information transmission protocols and to solicit information as to the information transmission protocol employed by the CSSP 71 to communicate with the information processing modules 43. The shelf service processor 74 can then be operable to respond with in-



formation defining the information transmission protocol employed by the CSSP 71 to communicate with the information processing modules 43. If the information transmission protocol used by the CSSP 71 is not one of those supported by the information processing module 43, the CSSP 71 could optionally indicate that the information processing module is incompatible with the CSSP 71. Alternatively, or in addition, the service controller 203 could make the determination as to whether it support the information transmission protocol identified by the shelf service processor 74 as that used by the CSSP 71 for communication with the information processing modules.

**[0210]** If the information transmission protocol employed by the CSSP 71 is supported by the information processing module (in this example an Ethernet protocol and an Infiniband protocol), the service controller 203 can then be operable to send an appropriate switching signal via the switching signal line (not shown) to the switching 440 and 442 to connect the information signal interface path including the appropriate information protocol interfaces (in this example the Ethernet protocol interface 207 or the Infiniband protocol interface 407). If the information transmission protocol employed by the CSSP 71 is not supported by the information processing module, then service controller 203 could prevent initialization of the information processing module from completing, possibly after reporting a fault via the management signal connection(s).

**[0211]** In the example shown in Figure 20B, the processor 192 is connected via respective signal paths 483 and 484 to first and second IPIs 207 and 407, respectively. The first IPI can be, in this example, an Ethernet protocol interface (EI) 207 and the second IPI can be, in this example, an Infiniband protocol interface (II) 407. The Ethernet protocol interface 207 can be connected via an information signal path 467 to a first input of switching logic 450 (e.g., multiplexer/demultiplexer logic). The Infiniband protocol interface 407 can be connected via an information signal path 468 to a second input of the switching logic 450 forming the switching logic. A common output from the switching logic 450 is connected via a common information signal path 491 to a common PLI 439, which can be implemented by a SERDES device as described earlier. The SERDES device 439 can be connected via a common information signal path 472 (in this case a PECL TX+/- and RX+/- signal line pair from the SERDES device 439) to an information connection 224 on the midplane connector 120. The switching logic 450 can be controlled via a common switching signal line (not shown) from, for example, the information processing module service controller (BSC 203) shown in Figure 11. The combination of the IPIs (e.g., Ethernet protocol interface 207 and Infiniband protocol interface 407), the switching logic (e.g., multiplexer/demultiplexer 450), and the common PLI (e.g., SERDES device 439), forms an information signal interface 431.

**[0212]** Further, in the example shown in Figure 20B, the processor 192 is also connected via respective signal paths 485 and 486 to third and fourth IPIs 208 and 408, respectively. The third IPI can be, in this example, an Ethernet protocol interface (EI) 208 and the fourth IPI can be, in this example, an Infiniband protocol interface (II) 408. The Ethernet protocol interface 208 can be connected via an information signal path 469 to a first input of switching logic 452 (e.g., multiplexer/demultiplexer logic). The Infiniband protocol interface 408 can be connected via an information signal path 470 to a second input of the switching logic 452 forming the switching logic. A common output from the switching logic 452 is connected via a common information signal path 492 to a common PLI 441, which can be implemented by a SERDES device as described earlier. The SERDES device 441 can be connected via a common information signal path 472 (in this case a PECL TX+/- and RX+/- signal line pair from the SERDES device 441) to an information connection 225 on the midplane connector 120. The switching logic 452 can also be controlled via a common switching signal line (not shown) from, for example, the information processing module service controller (BSC 203) shown in Figure 11. The combination of the IPIs (e.g., Ethernet protocol interface 208 and Infiniband protocol interface 408), the switching logic (e.g., multiplexer/demultiplexer 452), and the common PLI (e.g., SERDES device 441), forms an information signal interface 432.

**[0213]** As in the example described with reference to Figure 20A, on initialization, or during use, the information processing module service controller 203 can be operable to select one of the information protocol interfaces by sending a switching signal to the switching logic 450 and 452. The signal can be generated in response to user input at a management console connected though the hierarchical management structure described earlier, or in response to an automatic exchange of information between the service controller 203 in the information processing module 43 and the shelf service processor 74 in the CSSP 71 as described with reference to Figure 20A.

**[0214]** It will be appreciate that the example of Figure 20B enables the same functions to be performed as in the example of Figures 20A with the added advantages of saving costs associated with providing a separate PLI for each IPI.

**[0215]** It will be appreciated from the above that each of the examples of Figures 20A to 20B provide the advantage of enabling the automatic selection of one of a plurality of information protocol interfaces at initialization or during use, thereby enabling the economies of scale to be employed by the provision of information processing modules capable of operating under multiple information protocols. Although in these examples multiplexer/demultiplexers are employed as switching logic, it will be appreciated that other forms of switching logic could be employed in other examples.



[0216] The examples of each of Figures 19A, 19B, 20A and 20B can employ a compatibility check along the lines of that described with reference to Figure 20A.

[0217] Thus, for example, initialization software or firmware in the information processing module, for example in the PROM 201 shown in Figure 11, can be operable on initialization to cause the service controller 203 to send management information to the shelf service processor 47 to inform the shelf service processor 74 of the information transmission protocol or protocols it can support and/or to solicit information as to the information transmission protocol employed by the CSSP 71 to communicate with the information processing modules 43. The shelf service processor 74 can then be operable to indicate whether the information protocol or protocols that the information processing module 43 can support is, or are, compatible with the information transmission protocol employed the CSSP 71 to communicate with the information processing modules 43. If the service controller 203 receives a reply indicating that the information transmission protocol(s) supported by the information processing module 43 is, or are, not compatible with the information transmission protocol employed by the CSSP 71, the service controller 203 can be operable in response thereto to prevent initialization of the information processing module from completing, possibly after reporting a fault via the management signal connection(s). Otherwise the service controller 203 can be operable to complete initialization of the information processing module using the appropriate information transmission protocol.

[0218] Alternatively, or in addition, on initialization, or in use, the shelf service processor 74 can provide information to the service controller 203 in the information processing module identifying the information transmission protocol employed by the CSSP 71 to communicate with the information processing modules 43. The service controller 203 can then be operable to determine whether the information transmission protocol identified by the shelf service controller 74 is supported by the information processing module. If the identified information transmission protocol is not supported by the information processing module, then the service controller 203 can be operable to prevent initialization of the information processing module from completing, possibly after reporting a fault via the management signal connection(s). Otherwise the service controller 203 can be operable to complete initialization of the information processing module using the appropriate information transmission protocol.

[0219] Figures 21A and 21B are yet further examples in which multiple information protocol interfaces can be implemented following an initialization process as described above. In these examples, one or more field programmable gate arrays (FPGAs) are provided that can be programmable at initialization. These examples also give the advantage of enabling the automatic selection of one of a plurality of information protocol interfaces at

initialization or during use, thereby enabling the economies of scale to be employed by the provision of information processing modules capable of operating under multiple information protocols.

[0220] In the example shown in Figure 21A, the processor 192 is connected via respective signal paths 461 and 462 to first and second FPGA 560 and 562, respectively. The first FPGA 560 selectively implements, for example, the functions of an Ethernet protocol interface (EI), an Infiniband protocol interface (II), or any other appropriate information protocol interface 561 by downloading appropriate programming at initialization following an exchange of information as described above. An output from the FPGA 560 is connected via an information signal path 491 to a PLI 439, which can be implemented by a SERDES device as described earlier. The PLI 439 can be connected via an information signal path 472 (in this case a PECL TX+/- and RX+/- signal line pair from the SERDES device 439) to an information connection 224 on the midplane connector 120 as described earlier. The second FPGA 562 selectively implements, for example, the functions of an Ethernet protocol interface (EI), an Infiniband protocol interface (II), or any other appropriate information protocol interface 563 by downloading appropriate programming at initialization following an exchange of information as described above. An output from the FPGA 562 can be connected via an information signal path 492 to a PLI 441, which can be implemented by a SERDES device as described earlier. The PLI 441 can be connected via an information signal path 473 (in this case a PECL TX+/- and RX+/- signal line pair from the SERDES device 441) to an information connection 225 on the midplane connector 120 as described earlier. The FPGA could be programmable by the blade service controller 203 (see Figure 11) in response to determining the protocol supported by the shelf service processor 74, for example at initialization, with the necessary programming being downloaded from the shelf service processor 74 via the management connections 226, 227, or from storage provided in the blade 43. The programming could alternatively be controlled by the shelf service processor 74. As an alternative, the functions of the FPGAs 560 and 562 could be implemented in a common FPGA (not shown).

[0221] Figure 21B corresponds generally to Figure 21A, except that in this example, the FPGAs 660 and 662 implement additional functions. In particular, the FPGA 660 implements the functions of an information protocol interface 561 of the FPGA 560 plus that of the SERDES PLI 439 and the FPGA 662 implements the functions of an information protocol interface 563 of the FPGA 562 plus that of the SERDES PLI 441. The functions of the FPGAs 660 and 662 could be implemented in a common FPGA (not shown).

[0222] Figures 22A and 22B show examples of configurations in which components are saved and thereby costs can be reduced. In the example of Figure 22A, the

cost of an information protocol interface IPI (e.g., an Ethernet interface, a Infiniband interface, or the like) can be saved by selectively connecting a common IPI 439 to the SERDES PLIs 409 and 410 via switching logic (SL) 440. This embodiment permits savings by reducing the number of IPIs needed. Figure 22B shows another example in which a common IPI 439 and a common PLI 430 are selectively connected via switching logic (SL) 450 to one of two information connections. This example permits further savings by also reducing the number of PLIs needed. It is noted that the cost of each IPI is currently of the order of US\$40 and the cost of each SERDES is currently of the order of US\$5.

[0223] Turning firstly to Figure 22A, it will be noted that the processor 492 is connected to a common IPI 439 via the processor bus 498 and a bus connection 432. The common IPI 439 could be in the form of an Ethernet Interface, an Infiniband interface, or indeed any other appropriate information protocol interface. The IPI 439 is also connected to switching logic (SL) 440 via an information signal path 433. For example the switching logic 440 can be a 1:2/2:1 multiplexer/demultiplexer and can be controlled by the IPI 439 via line 436. The switching logic 440, for example a multiplexer/demultiplexer, provides respective information signal paths to first and second PLIs 409 and 410. The PLIs 409 and 410 are connected in turn to respective information connections 417 and 418 on the midplane connector 420 (which corresponds to the midplane connector 120 shown in Figure 11) via respective information signal paths 415 and 416 (e.g., respective PECL TX+/- and RX+/- signal line pairs). The combination of the common IPI (439), the switching logic 440 and the PLIs (409, 410) forms an information signal interface (431).

[0224] The IPI 439 can be operable on initialization to direct the switching logic to switch to a predetermined one of the PLIs 409/410 (say PLI 409), whereby the IPI 439 is then operable to communicate over the information connection 417 via the PLI 409. The IPI 439 can be operable to check the operation of the currently used connection (here via the information connection 417) on a regular basis, for example by "pinging" the switch in the CSSP 71 to which the information connection is connected. If the ping fails (that is no response is received within a predetermined time), then the IPI 439 can be operable to determine that this connection is faulty and switch the switching logic to the other of the PLIs 409/410 (say PLI 410). The IPI is then operable to communicate over the information connection 418 via the PLI 410. It will be appreciated that the embodiment described with reference to Figure 22A is able to provide redundant information connections while saving the cost of an IPI with respect to the arrangement illustrated, for example, in Figure 19A or Figure 19B.

[0225] Turning now to Figure 22B, it will be noted that the processor 492 is connected to a common IPI 439 via the processor bus 498 and a bus connection 432. The common IPI 439 could be in the form of an Ethernet

Interface, an Infiniband interface, or indeed any other appropriate information protocol interface. The IPI 439 is also connected to a common PLI 430 via an information signal path 433. The PLI 430 is connected in turn to switching logic 450, for example a multiplexer/demultiplexer, via information signal paths 443. The switching logic 450 can be, for example, a 2:4/4:2 multiplexer/demultiplexer, providing 1:2 multiplexing and 2:1 demultiplexing for each signal line of the differential PECL TX+/- and RX+/- signal line pair. In this example the switching logic 450 is controlled by the IPI 439 via line 436. The switching logic 450 provides respective information signal paths 445 and 446 (e.g., respective PECL TX+/- and RX+/- signal line pairs) to first and second information connections 417 and 418 on the midplane connector 420. The combination of the common IPI (407, 408), the PLIs (409, 410) and the switching logic 450 forms an information signal interface (431).

[0226] The IPI 439 can be operable on initialization to direct the switching logic to switch to a predetermined one of the information connections 417/418 (say information connection 417), whereby the IPI 439 is then operable to communicate over the information connection 417. The IPI 439 can be operable to check the operation of the currently used connection (here via the information connection 417) on a regular basis, for example by "pinging" the switch in the CSSP 71 to which the information connection is connected. If the ping fails (that is no response is received within a predetermined time), then the IPI 439 can be operable to determine that this connection is faulty and switch the switching logic to the other of the information connections 417/418 (say information connection 418). The IPI is then operable to communicate over the information connection 418. It will be appreciated that the embodiment described with reference to Figure 22B is able to provide redundant information connections while saving the cost of an IPI and a PLI with respect to the arrangements illustrated in Figures 19A and 19B, for example.

[0227] Although it would appear that a significant compromise is made by use of a common IPI and/or PLI in the embodiments of Figures 20A and 20B compared provision of two redundant IPIs and PLIs and in Figure 19, it has in practice been found that such components are extremely reliable, and the likelihood of failure of those components is negligible. In practice, it is typically the physical interconnections between the modules and/or software applications that are potentially unreliable. Accordingly, an embodiment of the invention is able to provide multiple redundant information connections while using shared components to reduce the cost without noticeably impacting reliability.

[0228] There has been described an information processing module for a modular computer system. The information processing module includes processor means for performing information processing, information connection means for communicating information signals, serializer/deserializer means connected be-

tween the processor means and the information connection means, information protocol interface means connected between the processor means and the serializer/deserializer means.

[0229] For example a blade server can be provided for a high density computer system. To enhance reliability, multiple redundant information connections can be provided from the server blade. In the specifically described examples of the invention, the use of a serializer/deserializer as a physical layer interface in a blade server means that the communication of information signals can be effected using any one of a plurality of different information protocols, by way of example only, an Ethernet or an Infiniband protocol. The decision on which information protocol to be used in a particular example can be decided in advance by providing a selected information protocol interface in the information processing modules. Alternatively, the selection of the appropriate information protocol interface can be effected dynamically at initialization or in use of the information processing module by including a plurality of information protocol interfaces for different protocols and enabling selection of the appropriate information protocol interface. Although in the described examples, two information protocols are described, namely an Ethernet protocol and an Infiniband protocol, it will be appreciated that the invention is not limited thereto and that other information protocols and/or more than two information protocols can be supported.

[0230] There has also been described an information processing module for a modular computer system. The information processing module includes processor means for performing information processing. A plurality of information connections means can be used to communicating information signals. Information signal interface means can connect the processor means to the information connection means. The information signal interface means can include common information protocol interface means and at least one physical layer interface means. As described above, the use of the common, or shared, information protocol means can reduce the cost of providing multiple information connections.

[0231] For example a blade server can be provided for a high density computer system. To enhance reliability, multiple redundant information connections can be provided from the server blade. To reduce the cost of providing such multiple redundant information connections, a common information protocol interface can be shared between information connections. A common physical layer interface can also be shared between information connections. For example, switching logic can be provided between the common physical layer interface and the information connections. Alternatively, or in addition, a common information protocol interface can be connected via switching logic to multiple physical layer interfaces. The multiple physical layer interfaces may each be provided one per information connection, or alternatively, they may each form common physical layer

interfaces connected via respective switching logic to multiple information connections.

[0232] There has further been described an information processing module for a modular computer system, the information processing module including processor means for performing information processing, module information connection means for communicating information signals, module information signal interface means connected between the processor means and the module information connection means, module management connection means for communicating management signals, and service controller means operable to perform system management functions, the service controller means being connected to the module management connection means. This configuration enables the provision of a processor technology independent information server blade for a modular server system. Thus, as described above, a server blade can be implemented using one or more processors selected, merely by way of example, from an UltraSPARC processor, an Intel x86 series and compatible processors, AMD x86 compatible processors, Alpha processors and PowerPC processors.

[0233] Although the embodiments above have been described in considerable detail, numerous variations and modifications will become apparent to those skilled in the art once the above disclosure is fully appreciated. It is intended that the following claims be interpreted to embrace all such variations and modifications.

## Claims

1. An information processing module for a modular computer system, the information processing module comprising at least one processor operable to perform information processing, at least one module information connection for communicating information signals, at least one serializer/deserializer connected between said at least one processor and said at least one module information connection, and at least one information protocol interface connected between said at least one processor and said at least one serializer/deserializer.
2. The information processing module of claim 1, comprising a plurality of information protocol interfaces for different information protocols.
3. The information processing module of claim 2, wherein each protocol interface is connected to said at least one module information connection via a respective serializer/deserializer.
4. The information processing module of claim 3, comprising switching logic selectively connecting said serializer/deserializers to said at least one module information connection, whereby an information

protocol interface for a selected information protocol is selectively connectable to at least one said module information connection via a respective serializer/deserializer and said switching logic.

5. The information processing module of claim 2, comprising switching logic selectively connecting a plurality of information protocol interfaces for different information protocols to a common serializer/deserializer, whereby an information protocol interface for a selected information protocol is selectively connectable at least one said module information connection via said switching logic and said common serializer/deserializer.
6. The information processing array of any of claims 2 to 5, comprising a plurality of information protocol interfaces for different information protocols implemented in one or more field programmable gate arrays.
7. The information processing module of any preceding claim, wherein at least one said information protocol interface is an Ethernet interface.
8. The information processing module of any preceding claim, wherein at least one said information protocol interface is an Infiniband interface.
9. The information processing module of any preceding claim, comprising a plurality of module information connections.
10. The information processing module of any preceding claim, comprising two module information connections, each operable to carry a selectable one of a plurality of information protocols.
11. The information processing module of any preceding claim, operable to determine whether it provides an information protocol interface compatible with the modular computer system, the information processing module operating with a compatible information protocol interface and terminating operation if there is no compatible information protocol interface.
12. The information processing module of any preceding claim, wherein at least one common module information protocol interface is connected to at least one said serializer/deserializer, said at least one common module information protocol interface being connectable to at least one selectable module information connection via said at least one serializer/deserializer.
13. The information processing module of claim 12, wherein said at least one common module informa-

tion protocol interface is connected to at least one common module serializer/deserializer, switching logic is connected between said at least one common module serializer/deserializer and said module information connections, and said switching logic is operable selectively to connect said at least one common module serializer/deserializer to at least one of said module information connections.

14. The information processing module of claim 13, comprising a single common module physical layer interface.
15. The information processing module of claim 12, comprising switching logic connected between said at least one common information protocol interface and a plurality of module serializer/deserializers, wherein said switching logic is operable selectively to connect said at least one common module information protocol interface to at least one of said plurality of module serializer/deserializers.
16. The information processing module of claim 15, comprising a single common module information protocol interface.
17. The information processing module of any of claims 13 to 16, wherein said switching logic is controlled by a signal generated by at least one said module information protocol interface.
18. The information processing module of any preceding claim, further comprising at least one module management connection for communicating management signals and at least one service controller operable to perform system management functions, said at least one service controller being connected to said at least one module management connection.
19. The information processing module of any preceding claim, further comprising an internal storage device.
20. The information processing module of any preceding claim, comprising an enclosure.
21. The information processing module of any preceding claim, configured as a field replaceable server blade.
22. A modular computer system comprising a carrier comprising at least one module receiving location configured to receive an information processing module according to any preceding claim, the module receiving location comprising at least one carrier information connection operable to interconnect with a corresponding module information connection.

tion of a received information processing module.

23. A modular computer system of claim 22, wherein the carrier comprises:

at least one switch connectable to at least one carrier information connection of at least one module receiving location and operable to connect the information processing module to an external information network in accordance with at least one information protocol; and at least one service processor operable to interface with a service controller of an information processing module to determine whether the information processing module comprises an information protocol interface compatible with an information protocol employed by the carrier.

24. The computer system of claim 23, wherein the carrier is configured to receive a plurality of information processing modules in respective module receiving locations:

25. The computer system of claim 24, wherein each switch is connectable to a carrier information connection at each module receiving location and is operable selectively to connect each received information processing module to an external information network.

26. The computer system of claim 25, wherein each information processing module comprises two module information connections, each module receiving location comprises two carrier information connections operable to interconnect with a respective module information connection of a received information processing module and the carrier includes two switches, each connectable to a respective carrier information connection of at least one module receiving location.

27. The computer system of claim 26, wherein each switch is connectable to a respective carrier information connection at each module receiving location.

28. The computer system of claim 26 or claim 27, comprising first and second switch modules removably mountable in the carrier, each switch module comprising a respective switch.

29. The computer system of claim 28, wherein each switch module further comprises a respective service processor to form a combined switch and service processor module.

30. The computer system of any one of claims 22 to 29,

comprising an interconnection member, the interconnection member carrying information connectors and interconnection pathways, said interconnection pathways being configured to support a plurality of information protocols:

31. The computer system of any one of claims 22 to 30, wherein at least one information protocol is an Ethernet protocol.

32. The computer system of any one of claims 22 to 31, wherein at least one information protocol is an Infiniband protocol.

33. The computer system of any one of claims 22 to 32, wherein the carrier is a rack mountable shelf.

34. The computer system of claim 26, configured as a high density server system, comprising at least one information processing module configured as a server blade.

35. A method of interconnecting an information processing module to other components of a modular computer system wherein the information processing module includes at least one processor operable to perform information processing, the method comprising:

selecting an appropriate information protocol interface; and  
said at least one processor communicating information signals via said appropriate information protocol interface and a serializer/deserializer to an information connection.

36. The method of claim 35, comprising communicating information signals via a selectable one of a plurality of information protocol interfaces supporting different information protocols.

37. The method of claim 35 or claim 36, wherein the information processing module is a server blade and the modular computer system is a high density server system.

38. The method of any of claims 35 to 37, wherein said information signals are passed from said information connection via a connection plane to other modules of the modular computer system.

39. The method of any of claims 35 to 38, comprising determining whether the information processing module supports an information protocol interface compatible with the modular computer system, and terminating operation if there is no compatible information protocol interface.

40. The method of any of claims 35 to 39, comprising performing a compatibility check as to whether an information processing module supports an information protocol employed for information communication between the information processing modules and other components of the modular computer system and selecting a compatible one of a plurality of information protocol interfaces provided by the information processing module.

5

10

15

20

25

30

35

40

45

50

55




Fig. 1




Fig. 2